Volume V: Unified Closure

Ahmed Ali Saleh

October 16, 2025

Contents

I Unified Field Closure Charter 1

Unified Field Closure Charter 3

1 Einstein's Dream, Operationally Framed 3

Einstein's Dream, Operationally Framed 3

1.1 Scope, Targets, and Acceptance Gates 4

1.1.1 Objectives 4

1.1.2 Milestones 6

1.1.3 Acceptance Gates (A1–A3) 9

1.2 Global Hypotheses and DoV 11

Global Hypotheses and DoV 11

1.2.1 DoV and Boundary Triggers 13

1.3 Conventions and Notation 14

Conventions and Notation 14

1.4 Acronyms and Shorthands 14

Acronyms and Shorthands 14

1.5 Core Symbol Inventory 15

Core Symbol Inventory 15

1.6 Proof Obligations and Artifact Policy 15

Proof Obligations and Artifact Policy 15

1.7 Reference Scales and Calibration Policy 15

Reference Scales and Calibration Policy 15

1.8 Crosswalk: Vols. I–IV to MSP Axioms 15

Crosswalk: Vols. I–IV to MSP Axioms 15

1.9 DoV and Boundary Triggers (Pointer) 16

DoV and Boundary Triggers (Pointer) 16

II Damascus Interpretation: Compact Proofs, Bounds, and

Closure 17

Damascus Interpretation: Compact Proofs, Bounds, and Closure 19

2 Hilbert Space from $\mathcal{L}^{(2)}$ and Pointer Sectors 19

Hilbert Space from $\mathcal{L}^{(2)}$ and Pointer Sectors 19

2.1 Pre-symplectic Reduction and GNS . . . 19
2.1.1 Green-hyperbolicity; causal propagators . . . 21
2.1.2 Peierls = CCR/CAR (BRST kernel at ghost #0) . . . 22
2.1.3 Hadamard parametrix on (M, g_{eff}) with fiber couplings . . . 24
2.2 Stabilization by large (ρ, η): Decoherence-free Sectors . . . 27
2.2.1 Spectral locking rates; Lyapunov exponents; gaps . . . 27
2.2.2 Cauchy-surface independence . . . 29
2.3 Observable Algebra, Superselection, and Coarse-Graining . . . 30
2.3.1 Fiber Noether charges . . . 30
2.3.2 Pointer POVMs and robustness margins . . . 31

3 Born's Rule from Spectral Flooding 35
3.1 Energy-share Dynamics: Balance Equations, Flooding Kernels, Martingales 35
3.1.1 Balance equations & kernels . . . 36
3.1.2 Martingales; concentration . . . 37
3.2 Stationary Limit: **Theorem (Born-limit)** and Hypotheses . . . 38
3.2.1 Hypotheses (gap/mixing, calibrated channels) . . . 38
3.2.2 Limit and identification $p_i = \mathrm{Tr}(\Pi_i \rho_\infty)$ . . . 39
3.3 Deviation Bounds for Finite Locking Rates . . . 41
3.3.1 Sub-Gaussian finite-time bounds . . . 41
3.3.2 Error vs. γ scaling . . . 42

4 GKSL in the Spectral Medium: Locality, CP, and No-signalling 45
4.1 Microscopic Derivation: System–Meter Coupling, CP Semigroups, Instruments . . . 45
4.1.1 Derivation under locality radius L_{loc} . . . 45
4.1.2 Instrument maps and POVMs . . . 47
4.2 **Theorem (No-signalling):** Local GKSL $\Rightarrow$ No Superluminal Signalling 48
4.2.1 Spacelike separation on (M, g_{eff}) . . . 48
4.2.2 Colored noise via CP time-local embedding . . . 49
4.3 POVMs and Spectral Error–Disturbance (Ozawa/Branciard) with Colored Corrections . . . 50

5 Colored Noise, Zeno Engineering, and Thermodynamic Ledgers 55
5.1 Non-Markovian Kernels → Time-local Embeddings (Pseudomodes; CP Tests) 55
5.1.1 Embedding constructions . . . 56
5.1.2 CP diagnostics . . . 57
5.2 Colored-Zeno Control: Filter Functions, Convex LMIs, Robustness Margins 58
5.2.1 Design of filter functions . . . 59
5.2.2 LMI robustness . . . 60
5.3 Work/Heat Ledgers: First/Second-law Balance; Landauer-type Bounds . 62
5.3.1 Ledger construction . . . 63
5.3.2 Bounds and consistency . . . 64

III Entanglement, Threads, and Device-Independent Constraints 67

6 χ^{μ}-Threads and Structural Nonseparability 69
6.1 Thread Graphs, Cuts, and Monogamy Inequalities . . . 69
6.1.1 Graph models and cuts . . . 69
6.1.2 Monogamy relations . . . 71
6.2 Thread-adjusted Measures (Negativity/Rel. Entropy): Continuity, Faithfulness . . . 72
6.2.1 Adjusted negativity . . . 72
6.2.2 Relative-entropy based . . . 74
6.3 Witnesses and Tomography on g_{eff} Backgrounds; Finite-speed Effects . . 75
6.3.1 SDP witnesses . . . 75
6.3.2 Tomography under cone constraints . . . 77

7 Tsirelson Saturation without Superluminal Signalling 79
7.1 Operator-norm CHSH; SDP Bounds; Saturating States/Measurements . . 79
7.1.1 CHSH via operator norms . . . 79
7.1.2 SDP bounds & saturators . . . 81
7.2 Structure Proof: GKSL Locality ⇒ No-signalling beyond Markovian Embeddings . . . 83
7.2.1 Beyond-Markov embeddings . . . 84
7.3 Device-independent Feasibility Regions under Colored Noise; Finite Statistics 86
7.3.1 Concentration bounds . . . 86

IV Master Spectral Principle and Unified Completion 89

8 Master Spectral Principle (MSP): Minimal Axioms and Corollaries 91

8.1 Minimal Axiom Set Compressing Vols. I–IV; Independence and Redundancy Checks . . . 91
8.1.1 S0–S4 (minimal) . . . 91
8.1.2 Independence/redundancy table . . . 92
8.2 Corollaries: Emergent g_{eff} and Ehresmann→YM Connections; Anomalies/Inflow . . . 94
8.2.1 Emergent cone/connection . . . 94
8.2.2 Anomalies and inflow . . . 96
8.3 Ward/BRST Structures as MSP Consequences; Physical-Quotient Equivalence 97
8.3.1 ST identities . . . 97
8.3.2 Physical-quotient equivalence . . . 98

9 Geometric Determination of Couplings and Mixings 101
9.1 Lie-metric on Iso($\mathcal{F}$) and $g_{1,2,3}$ Relations; Thresholds; Positivity . . . 101
9.1.1 Metric relations . . . 101
9.1.2 Threshold/positivity . . . 103
9.2 Weinberg Angle from Fiber Geometry: **Prediction with Uncertainties** 104
9.2.1 Reference-scale prediction . . . 104
9.2.2 RG transport to M_Z . . . 106
9.3 Overlap-driven Yukawa Textures; CKM/PMNS Sum Rules . . . 107
9.3.1 Overlap textures . . . 107
9.3.2 CKM/PMNS relations . . . 109

V Predictive Closure and Kill Signatures 113

10 Compact Master Predictions 115
10.1 Gauge-coupling Ratios and $\sin^2\theta_W$ at a Reference Scale; RG to M_Z . . . 115
10.1.1 Reference-scale envelopes . . . 115
10.1.2 RG maps and uncertainties . . . 116
10.2 Precision Baselines: (S, T, U) and EDM Floors . . . 117
10.2.1 Oblique parameters . . . 118
10.2.2 Irreducible EDM floors . . . 119
10.3 Neutrino NSI/CPT-odd Signatures and Appearance Channels . . . 121
10.3.1 NSI matrix patterns . . . 121
10.3.2 Waveform shapes & baselines . . . 123

11 Falsifiability and Acceptance Windows 125
11.1 Positivity/Hyperbolicity Gates; Cone-LMI Feasibility Regions . . . 125
11.1.1 Forward-limit constraints . . . 125

11.1.2 Cone LMIs 128
11.2 Near–Far Detector Consistency; Orthogonal Templates to Break Degeneracies 130
11.2.1 ND–FD maps 130
11.2.2 Orthogonal templates 131
11.3 Decision Tables: Which Observable Kills Which Hypothesis; Revision Paths 133
11.3.1 Kill windows 134
11.3.2 Revision playbook 135

VI Tabletop Quantum-Control Probes 137

12 Colored-Zeno Protocols and Noise Spectroscopy 139
12.1 Filter-function Design and Robustness (Convex LMIs); Reference Sequences/Code 139
12.1.1 Design toolchain 142
12.1.2 Robustness margins 143
12.2 Parameter-to-Knob Maps: Cold Atoms, Waveguides, (Optional) Superconducting Circuits 144
12.2.1 Cold-atom implementations 145
12.2.2 Waveguide implementations 146
12.3 Sensitivity Forecasts and Systematics Ledgers; Stopping Rules 148
12.3.1 Forecast methodology 148
12.3.2 Systematics ledger & stopping 149

13 Implementations with Cold Atoms and Waveguides 153
13.1 Experimental Sequences and Readouts; Calibration/Health Checks 153
13.1.1 Sequences & readouts 153
13.1.2 Calibration & checks 155
13.2 Verification and Tomography (GST/RB under Colored Noise); Error Budgets 156
13.2.1 GST/RB under color 156
13.2.2 Error budgets 158
13.3 Bridging to MSP Predictions: Maps to (S, T, U), EDM Floors, NSI Projections; Release Artifacts 159
13.3.1 Bridge maps 160
13.3.2 Release artifacts 161

Epilogue: Unified Closure Statement 165

What Has Been Proved and What Can Kill It 165

Conceptual Afterword: Event-Waves, Spectral Flooding, and Emergence **171**

Appendix: Concept–Mathematics Bridge **175**

Part I

Unified Field Closure Charter

Chapter 1

Einstein's Dream, Operationally Framed

Prelude. This volume adopts an *operational* stance: only quantities with units, device bridges, and a declared *domain of validity* (DoV) are admitted. The spectral medium is described by invariant, unit-bearing descriptors $\mathcal{S} = (\rho, \eta, \phi, T, \chi^\mu)$ on an emergent background (M, g_{eff}). Fiber data (F, h) capture internal geometry; local isometries of (F, h) give rise to Yang–Mills connections via Ehresmann transport. The *Master Spectral Principle* (MSP) compresses Vols. I–IV into a finite set of axioms and gates ensuring stability (S), causal hyperbolicity (G), complete positivity (TQ), and nonnegative dissipation (D). "Unified closure" means two things: (i) the quantum foundations are completed by a *Born-limit* and a *no-signalling* theorem within the spectral medium; (ii) the high-energy sectors are *geometrically* tied to a small sovereign input set, delivering numerical, falsifiable predictions with explicit uncertainty budgets.

Notation. We use $\hbar$ and c explicitly where clarity benefits; Tr is the trace, Spec the spectrum; projectors are Π_i. The Liouvillian (GKSL) generator is $\mathcal{L}$. The locality radius on (M, g_{eff}) is L_{loc}. The Liouvillian mixing gap on pointer sectors is $\Delta_{\mathcal{L}} > 0$. Geometric gauge normalization is κ_G; counting factors for matter families are $(\mathcal{N}_u, \mathcal{N}_d, \mathcal{N}_\ell)$.

Definition 1.1 (Domain of Validity (DoV))**.** The *Domain of Validity* is the maximal connected subset $\mathcal{D} \subset \mathcal{P}$ of the parameter space (scales, densities, temperatures, coupling strengths, bandwidths, geometry) such that, on $\mathcal{D}$:

(i) the cone/causality LMIs and forward-limit positivity constraints hold (Sections 11.1.1 and 11.1.2);

(ii) a CP time-local (possibly colored) embedding exists for the dynamics (Chapter 4, §4.2.2);

(iii) the Liouvillian mixing gap on pointer sectors satisfies $\Delta_{\mathcal{L}} > 0$ (Section 3.3);

(iv) calibration tolerances and DI/cone slacks lie within declared budgets.

Boundary triggers: gap closure $\Delta_{\mathcal{L}} \downarrow 0$, loss of CP-divisibility, violation of locality radius L_{loc}, or any negative eigenvalue in the cone/forward LMIs.
Operational procedures to establish and audit DoV *are detailed in Vol. VI, Sec. §⟨DoV-audit chapter ref⟩.*

1.1 Scope, Targets, and Acceptance Gates

This section records the *targets* of the volume and the *gates of acceptance.* Targets are phrased as mathematical statements to be proved in Parts I–III and as calibrated prediction envelopes verified in Parts IV–V. Acceptance gates are binary checks—logical and numerical—that a result must pass to be part of the final closure set.

1.1.1 Objectives

Unified closure.

The objective is to close the theory by proving, under explicit hypotheses (H0–H5), both a *Born-limit* and a *no-signalling* theorem in the spectral medium, and by compressing Vols. I–IV into a minimal axiom system we call the *Master Spectral Principle* (MSP), from which the gauge sectors and their precision phenomenology follow.

Definition 1.2 (Master Spectral Principle (MSP))**.** An MSP-instance is a quintuple

$$\mathfrak{M} = \Big(M, g_{\text{eff}};\ F, h;\ S_{\text{master}};\ \text{Gates}\Big),$$

where: (i) D_μ is an Ehresmann derivative on the fiber bundle $E \to M$ with fiber (F, h); local isometries $\text{Iso}(F, h)$ induce a Yang–Mills connection A_μ with curvature $\mathcal{F}_{\mu\nu}$; (ii) S_{master} is a spectral action on fields $\mathcal{S}$ and gauge data $(A_\mu, \mathcal{F})$, compatible with g_{eff} and device bridges; (iii) the *gates* $\{\mathsf{S}, \mathsf{G}, \mathsf{TQ}, \mathsf{D}\}$ hold:

$$\begin{aligned}
\mathsf{S}: &\quad \text{positivity of the quadratic form and absence of ghosts,}\\
\mathsf{G}: &\quad \text{hyperbolic principal symbol; forward cones of } g_{\text{eff}},\\
\mathsf{TQ}: &\quad \text{complete positivity of reduced dynamics (GKSL or CP embeddings),}\\
\mathsf{D}: &\quad \dot{S}_{\text{vN}} \geq 0 \text{ for admissible instruments.}
\end{aligned}$$

Two structural theorems formalize "unified closure":

Theorem 1.3 (Born-limit (structural form))**.** *Let H0–H5 hold, with $\Delta_{\mathcal{L}} > 0$ on calibrated pointer sectors $\{\Pi_i\}$ and $\mathcal{L}$ a local GKSL generator on (M, g_{eff}). Then the energy-share dynamics admit a stationary limit $\rho_\infty = \lim_{t\to\infty} e^{t\mathcal{L}}\rho_0$ and*

$$p_i \;=\; \text{Tr}(\Pi_i\, \rho_\infty)\,.$$

Moreover, for finite locking rate γ_{lock} there exist sub-Gaussian deviation bounds

$$\mathbb{P}\Big(\Big|\tfrac{1}{t}\int_0^t \mathrm{Tr}(\Pi_i\rho_s)\,\mathrm{d}s \;-\; p_i\Big| > \epsilon\Big) \;\le\; 2\exp\Big[-c\,\epsilon^2\,\gamma_{\text{lock}}\,t\Big]$$

for some constant $c > 0$ determined by the instrument's noise spectrum.

Theorem 1.4 (Local GKSL $\Rightarrow$ No-signalling)**.** *Let regions $X, Y \subset M$ be spacelike separated with respect to g_{eff}, and let $\Lambda_t^{X\cup Y}$ be the CP, trace-preserving evolution generated by a local Liouvillian $\mathcal{L} = \mathcal{L}_X \oplus \mathcal{L}_Y$ with support within radius L_{loc} and CP time-local embeddings for colored noise. For any initial ρ_{XY} and any instrument Instr_Y acting in Y,*

$$\mathrm{Tr}_Y\Big[\Lambda_t^{X\cup Y}\Big((\mathbb{I}_X \otimes \mathsf{Instr}_Y)(\rho_{XY})\Big)\Big] \;=\; \Lambda_t^X\Big(\mathrm{Tr}_Y\,\rho_{XY}\Big)\,,$$

i.e. marginal statistics in X are independent of operations performed in Y.

Single-sovereign inputs (MSP*)**.**

We aim to *minimize* the set of free inputs to a small sovereign tuple

$$\Theta_{\text{geom}} = \Big(\kappa_G, \mathcal{N}_u, \mathcal{N}_d, \mathcal{N}_\ell;\ \text{topological indices}\Big)$$

from which gauge couplings, mixing angles, and Yukawa textures are computed. Concretely, at a reference scale μ_0,

$$g_i^{-2}(\mu_0) = \kappa_G\ \mathrm{tr}_h(T_iT_i)\,\mathcal{N}_i \;+\; \delta_i[h,F]\,, \qquad i \in \{1,2,3\}, \tag{1.1}$$

$$\sin^2\!\theta_W(\mu_0) = f_{\text{geom}}\Big[h, \kappa_G; \mathcal{N}_u, \mathcal{N}_d, \mathcal{N}_\ell\Big] \;\pm\; \Delta_{\text{geom}}\,, \tag{1.2}$$

with δ_i encoding calculable thresholds and f_{geom} a fiber-geometric functional determined by (F, h). Transport to M_Z uses RG maps with an uncertainty budget Δ_{RG}.

Definition 1.5 (Sovereign minimization problem)**.** Let $\mathcal{C}_{\mathsf{SGTQD}}$ be the feasibility set of $(h, \kappa_G, \mathcal{N}_u, \mathcal{N}_d, \mathcal{N}_\ell)$ satisfying all gates. Given experimental targets $\widehat{g}_i(\mu_0)$, $\widehat{\sin^2\!\theta_W}(\mu_0)$ with covariances Σ_g, Σ_W, choose

$$(h^\star, \kappa_G^\star, \mathcal{N}_u^\star, \mathcal{N}_d^\star, \mathcal{N}_\ell^\star) \;=\; \arg\min_{\mathcal{C}_{\mathsf{SGTQD}}} \Big[\Big(g(\mu_0)-\widehat{g}\Big)^\top \Sigma_g^{-1}\Big(g(\mu_0)-\widehat{g}\Big) + \frac{\Big(\sin^2\!\theta_W(\mu_0) - \widehat{\sin^2\!\theta_W}\Big)^2}{\Sigma_W}\Big],$$

and propagate to M_Z with RG to produce prediction envelopes.

Remark 1.6 (Acceptance gates A1–A3). A result is accepted into the closure set iff: **A1** it is stated as a formal Lemma/Theorem with explicit hypotheses; **A2** its artifacts are reproducible (cards/seeds/HDF5, scripts) and pass regression tests; **A3** it satisfies all gates $\mathsf{S}, \mathsf{G}, \mathsf{TQ}, \mathsf{D}$ within the declared DoV and triggers end-of-validity when violated.

The remainder of Part II turns Theorems 1.3 and 4.3 into full proofs under H0–H5, and Part IV instantiates Theorems 1.2 and 1.5 to derive the geometric relations (1.1)–(1.2) with quantified uncertainties.

1.1.2 Milestones

We track six milestones (M1–M6). Each is stated as a mathematical target with explicit hypotheses and an acceptance checklist. Cross–references point to the chapters where the formal statements and proofs/derivations appear.

M1 — Hilbert/Pointer.

Construct the quantum kinematics from the quadratic spectral action $\mathcal{L}^{(2)}$ on (M, g_{eff}), and define calibrated pointer sectors.

- **Kinematics.** Under Green–hyperbolicity of the kinetic operator K on (M, g_{eff}), the causal propagators $\Delta^{\pm}$ exist and the Peierls bracket equals the canonical (anti)commutator on the BRST physical quotient:

$$[\hat{\Phi}(f), \hat{\Phi}(g)] \;=\; i\,\Delta(f,g)\mathbb{I}, \qquad \Delta := \Delta^{-} - \Delta^{+},$$

  yielding a positive CCR/CAR algebra and a Hadamard GNS representation.

- **Pointer sectors.** Calibrated instruments/POVMs $\{\Pi_i\}$ define pointer sectors with Liouvillian mixing gap $\Delta_{\mathcal{L}} > 0$. Expectation values of pointer observables are Cauchy–surface independent within the declared DoV.

- *Acceptance:* (A1) Proposition/Theorem "Green $\Rightarrow$ CCR/CAR" and "surface independence"; (A2) reproducible seeds/HDF5 for propagators and Hadamard tests; (A3) S/Gsatisfied.

Lemma 1.7 (Peierls $\Rightarrow$ CCR/CAR on the BRST physical quotient)**.** *Let $\mathcal{L}^{(2)}$ be a quadratic Lagrangian on a globally hyperbolic (M, g_{eff}) with linear Euler–Lagrange operator P (bosonic or fermionic) that is Green-hyperbolic on the declared DoV. Denote by $G_{\text{ret/adv}}$ the retarded/advanced fundamental solutions and by $\Delta := G_{\text{ret}} - G_{\text{adv}}$ the causal propagator. Let $(\mathcal{F}, s)$ be the BRST complex with gauge-fixing so that the physical smearing space is $\mathcal{D}_{\text{phys}} := \ker P / \big(\operatorname{im} P + \operatorname{im} s\big)$. Define the Peierls bracket on representatives f, g by*

$$\{F_f, F_g\}_{\text{Peierls}} \;:=\; \int_M f\,\Delta\,g \;\equiv\; \sigma([f],[g]).$$

Then the $*$-algebra $\mathfrak{A}$ generated by smeared fields obeys, on the BRST physical quotient,

$$[\hat{\Phi}([f]), \hat{\Phi}([g])] \;=\; i\,\sigma([f],[g])\,\mathbb{I} \quad \textit{(bosons)}, \qquad \{\hat{\Psi}([f]), \hat{\Psi}([g])\} \;=\; \sigma([f],[g])\,\mathbb{I} \quad \textit{(fermions)}. \tag{1.3}$$

Sketch. *Green-hyperbolicity gives well-defined* $G_{\mathrm{ret/adv}}$ *and antisymmetric* Δ. *BRST reduction quotients gauge/constraint degeneracies so* σ *descends nondegenerately to* $\mathcal{D}_{\mathrm{phys}}$. *By Peierls' construction,* $\{F_f, F_g\}_{\mathrm{Peierls}}$ *coincides with the classical (pre-)symplectic form; canonical quantization then yields* (1.3). *Gates* S/G*ensure hyperbolicity/locality on the* DoV.

Proposition 1.8 (Hadamard ⇒ positive GNS representation)**.** *prop:hadamard-gns Assume a quasi-free state* ω *on* $\mathfrak{A}$ *whose two-point distribution* Λ_ω *on* (M, g_{eff}) *satisfies the microlocal Hadamard condition on the* DoV, *namely*

$$\mathrm{WF}(\Lambda_\omega) \;=\; \mathcal{C}^+ \subset T^*M \times T^*M, \qquad \Lambda_\omega - \tfrac{i}{2}\Delta \text{ is symmetric,}$$

with Δ *the causal propagator on the BRST-physical quotient. Then:*

(i) ω *is positive and defines a GNS triple* $(\mathcal{H}_\omega, \pi_\omega, \Omega_\omega)$*;*

(ii) π_ω *is a CCR/CAR representation with microcausality;*

(iii) the two-point function has Hadamard short-distance form $\Lambda_\omega = H + W$, *where* H *is the local parametrix and* $W \in C^\infty$.

Sketch. *The Hadamard wavefront set implies positivity of the two-point form on test data and ensures* $\operatorname{Im}\Lambda_\omega = \frac{1}{2}\Delta$, *so that the canonical commutation/anticommutation relations hold in* π_ω. *The GNS construction then yields a Hilbert space with cyclic vector* Ω_ω. A2 check. *On the* DoV, *certify Hadamard by testing that* $\Lambda_\omega - H$ *is smooth (finite-difference fits of the parametrix) and by verifying* $\mathrm{WF}(\Lambda_\omega) \subset \mathcal{C}^+$ *numerically on sampled charts.*

Proposition 1.9 (Cauchy-surface independence)**.** *prop:cauchy-indep Let* u, v *be smooth solutions of* $P\phi = 0$ *with compactly overlapping supports between two Cauchy surfaces* $\Sigma_1, \Sigma_2 \subset M$. *Let* $j^\mu(u, v)$ *be the conserved current associated with* $\mathcal{L}^{(2)}$ *(e.g. for a scalar,* $j^\mu = u\,\nabla^\mu v - v\,\nabla^\mu u$*). Then the (pre-)symplectic form*

$$\sigma_\Sigma(u, v) \;:=\; \int_\Sigma n_\mu\, j^\mu(u, v)\, \mathrm{d}\Sigma$$

is independent of the Cauchy surface: $\sigma_{\Sigma_1}(u, v) = \sigma_{\Sigma_2}(u, v)$. Sketch. *By the equations of motion,* $\nabla_\mu j^\mu = 0$. *Apply Stokes' theorem to the spacetime slab bounded by* $\Sigma_1 \cup \Sigma_2$*; boundary terms at spatial infinity vanish on the* DoV*(cone/growth bounds). Hence the flux of* j *through* Σ *is conserved, and the induced form—and thus expectation values in Hadamard GNS states—are Cauchy-surface independent.*

Acceptance 1.10. A1–A3 hooks for M1acc:M1-hooks A1: Theorem 1.7 and **????** provide formal statements.
A2: Publish HDF5 seeds with sampled Λ_ω, parametrix fits H, and numerical checks of WF containment; include CCR/CAR residuals.

A3: Gate S(hyperbolicity/Hadamard) and G(locality radius) logs accompany each certification on the DoV.

M2 — Born-limit.

Prove the stationary *Born-limit* for spectral flooding under H0–H5 with $\Delta_{\mathcal{L}} > 0$: there exists ρ_∞ such that $p_i = \mathrm{Tr}(\Pi_i\, \rho_\infty)$, with finite–time deviation bounds controlled by the locking rate γ_{lock}. (See Theorem 1.3 and §3.)

- *Acceptance:* (A1) Theorem with hypotheses and bounds; (A2) Monte–Carlo notebooks and variance budgets; (A3) TQ/Dsatisfied for the calibrated instruments.

M3 — Local GKSL ⇒ No-signalling.

Derive the microscopic GKSL generator under a locality radius L_{loc} and prove that local CP (time–local) dynamics preserve spacelike marginal independence:

$$\mathrm{Tr}_Y\Big[\Lambda_t^{X\cup Y}((\mathbb{I}_X \otimes \mathsf{Instr}_Y)(\rho_{XY}))\Big] \;=\; \Lambda_t^X(\mathrm{Tr}_Y\, \rho_{XY}).$$

(See Theorem 4.3 and §4.)

- *Acceptance:* (A1) Theorem and CP–embedding lemma for colored noise; (A2) channel–tomography scripts; (A3) G/TQsatisfied.

M4 — MSP ⇒ BRST/Ward.

Show that MSP axioms imply the standard gauge identities on the fiber bundle:

$$\begin{aligned} &\text{Ward:} \quad D_\mu \frac{\delta S_{\mathrm{master}}}{\delta A_\mu^I} + \xi_I^A(y)\, \frac{\delta S_{\mathrm{master}}}{\delta y^A} \;=\; 0, \\ &\text{BRST:} \quad s^2 = 0, \qquad \mathcal{S}(\Gamma) = 0 \text{ (Slavnov–Taylor).} \end{aligned} \tag{1.4}$$

Prove *physical–quotient equivalence*: observables modulo BRST–exact insertions coincide with the gauge–invariant quotient defined by (1.4). (See §8.)

- *Acceptance:* (A1) Lemma–chain: Ehresmann→YM, Ward, nilpotent BRST, ST; (A2) cohomology notes/artifacts; (A3) S/TQsatisfied.

M5 — Couplings/angles/textures.

Determine gauge couplings and the Weinberg angle from fiber geometry at a reference scale μ_0, and produce overlap–driven Yukawa textures:

$$g_i^{-2}(\mu_0) = \kappa_G\, \mathrm{tr}_h(T_i T_i)\, \mathcal{N}_i + \delta_i[h, F], \quad i = 1, 2, 3, \tag{1.5}$$

$$\sin^2\theta_W(\mu_0) = f_{\mathrm{geom}}\Big[h, \kappa_G; \mathcal{N}_u, \mathcal{N}_d, \mathcal{N}_\ell\Big] \;\pm\; \Delta_{\mathrm{geom}}, \tag{1.6}$$

$$(Y_f)_{ij} \propto \int_F \Big\langle \overline{\Psi}_i^{(f)},\, \Upsilon^{(f)}\, \Psi_j^{(f)} \Big\rangle_h\, \mathrm{dvol}_h, \qquad f \in \{u, d, \ell, \nu\}. \tag{1.7}$$

Diagonalization yields $V_{\mathrm{CKM}}, U_{\mathrm{PMNS}}$ with sum–rules reported in Part III.

- *Acceptance:* (A1) derivations of (1.5)–(1.7); (A2) RG notebooks and uncertainty budgets ($\Delta_{\text{geom}}, \Delta_{\text{RG}}$); (A3) S/Gsatisfied (positivity/hyperbolicity).

M6 — Predictions & Kill.

Publish compact prediction envelopes and *kill windows*: gauge–coupling ratios and $\sin^2\theta_W$ at μ_0 (RG–transported to M_Z), precision baselines (S, T, U), irreducible EDM floors, and neutrino NSI patterns. Define decision tables mapping each observable to the hypothesis it kills if outside its acceptance interval.

- *Acceptance:* (A1) decision tables with formulas and bounds; (A2) released artifacts (cards/seeds/HDF5) and CI regression tests; (A3) all gates hold within the declared DoV.

1.1.3 Acceptance Gates (A1–A3)

We adopt three *acceptance gates*. A result joins the closure set only if it satisfies **A1**–**A3**. Each gate is operational, with precise mathematical checks and end–of–validity triggers.

A1 — Formal Lemma/Theorem statements.

Every key claim is stated as a numbered Lemma/Theorem with: (i) explicit hypotheses (referencing H0–H5 and the declared DoV); (ii) a proof or a forward reference to a proof within this volume; (iii) boundary triggers (e.g. gap closure $\Delta_{\mathcal{L}}\downarrow 0$, loss of CP or locality) that invalidate the claim outside the DoV.

A2 — Reproducible artifacts.

Each result ships with a minimal, public artifact set: cards (human–readable summaries), seeds/HDF5 (arrays with shapes and digests), and notebooks/scripts sufficient to reproduce all figures and bounds. CI regression tests are provided to guard against drift.

A3 — Consistency with S/G/TQ/D gates.

The claim must satisfy the stability, causality, complete–positivity, and dissipation gates below, within the declared DoV.

Gate 1.1: S (Stability / Spectral Positivity)

Let $\mathcal{L}^{(2)}[\Phi] = \frac{1}{2}\int \mathrm{d}^4x \left(\partial_\mu \Phi^a K^{\mu\nu}_{ab} \partial_\nu \Phi^b + \Phi^a M_{ab} \Phi^b\right)$ be the quadratic spectral action on (M, g_{eff}). Stability holds if:

S1. $K^{\mu\nu}_{ab}$ defines a positive quadratic form on the physical (BRST) quotient: for all compactly supported test fields f, $\int f_a K^{00}_{ab} f_b \, \mathrm{d}^3x \geq 0$ and no ghost poles occur.

S2. The Källén–Lehmann spectral densities are positive: $\rho_{ab}(\mu^2) \succeq 0$, ensuring reflection positivity of 2–point functions and absence of tachyons within the

DoV.

Gate 1.2: G (Causality / Hyperbolicity)

Let $P(x,\xi) = Z^{\mu\nu}(x)\,\xi_\mu\xi_\nu$ be the principal symbol of the kinetic operator pulled back to (M, g_{eff}).

G1. (Hyperbolicity) For each $x \in M$ there exists a covector τ with $P(x,\tau) < 0$ such that the polynomial $P(x, \tau + \lambda\,\zeta)$ has only real roots in λ for all ζ. Equivalently, the characteristic cones of P are contained in the forward cones of g_{eff}.

G2. (Front–speed bound) Characteristic speeds obey $c_{\text{char}} \le c$, and causal propagators satisfy $\operatorname{supp}\Delta^{\pm} \subset J^{\pm}_{g_{\text{eff}}}$.

Violation of **G1**–**G2** triggers end–of–validity.

Gate 1.3: TQ (Quantum CP / Channel Consistency)

Reduced dynamics on calibrated pointer sectors are CP and (when applicable) trace–preserving:

TQ1. (GKSL form) In the Markovian window,

$$\dot\rho = -i[H,\rho] + \sum_{\alpha,\beta} C_{\alpha\beta}\Big(L_\alpha \rho L_\beta^\dagger - \tfrac{1}{2}\{L_\beta^\dagger L_\alpha, \rho\}\Big),$$

with Kossakowski matrix $C \succeq 0$.

TQ2. (Colored noise) Outside Markovianity, a time–local CP embedding exists: $\dot\rho_t = \mathcal{L}_t[\rho_t]$, and the associated dynamical map Λ_t is CP–divisible on the DoV.

TQ3. (Instrument consistency) Instruments are CP and obey data–processing inequalities; composition with Λ_t preserves CP.

Any negative eigenvalue in C or failure of CP–divisibility triggers end–of–validity.

Gate 1.4: D (Dissipation / Entropy Production)

Let ρ_t evolve under the admissible dynamics and ρ_{ss} be a steady (or instantaneous reference) state. Entropy production is nonnegative:

D1. (Spohn inequality) $\sigma(t) = -\dfrac{\mathrm{d}}{\mathrm{d}t}\,S(\rho_t\|\rho_{\text{ss}}) \;\ge 0$.

D2. (Thermodynamic ledger) For calibrated devices, $\dfrac{\mathrm{d}}{\mathrm{d}t}S_{\text{vN}}(\rho_t) - \beta\,\dot Q_t \ge 0$, and Landauer–type bounds hold for admissible erasure processes.

A negative $\sigma(t)$ within the DoV disqualifies the result.

Checklist 1.1: Acceptance A1–A3

A. **A1** Formal statement with hypotheses H0–H5, declared DoV, and boundary triggers.

B. **A2** Public artifacts: (i) card (human–readable); (ii) seeds/HDF5 with digests; (iii) notebooks/scripts; (iv) CI regression test.

C. **A3** Gate compliance: pass Section 1.1.3 on the same DoV.

1.2 Global Hypotheses and DoV

We work under six global hypotheses H0–H5. They pin down the kinematics of the spectral medium, how locality and noise are implemented, and how devices are calibrated. The hypotheses are deliberately *operational*: every symbol has units and an associated measurement protocol. Wherever possible we give the mathematical form used later in Parts I–III.

H0 (Spectral state and background).

The medium is specified by unit-bearing descriptors $\mathcal{S} = (\rho, \eta, \phi, T, \chi^\mu)$ pulled back along a section $s : M \to E$ of a fiber bundle $E \to M$ with fiber (F, h). The emergent background carries an effective metric g_{eff} and an Ehresmann derivative D_μ acting on fiber fields y^A:

$$D_\mu y^A = \partial_\mu y^A + \Gamma^A{}_B(y)\, \partial_\mu y^B - \xi^A_I(y)\, A^I_\mu, \qquad \mathcal{F}^I_{\mu\nu} = \partial_\mu A^I_\nu - \partial_\nu A^I_\mu + f^I{}_{JK} A^J_\mu A^K_\nu .$$

All constitutive relations and actions are $\mathrm{Iso}(F, h)$-covariant and respect the causal cones of g_{eff}.

H1 (Locality radius on (M, g_{eff})).

System–meter and matter–gauge couplings obey a finite locality radius L_{loc}: for any spacetime kernel $K(x, y)$ entering the microscopic interaction,

$$\operatorname{supp} K \subset \{(x, y) \in M^2 : \; d_{g_{\text{eff}}}(x, y) \le L_{\text{loc}}\} \quad \text{or} \quad |K(x, y)| \le C\, e^{-d_{g_{\text{eff}}}(x,y)/L_{\text{loc}}}.$$

Equivalently, admissible operations have support contained in (or exponentially concentrated inside) the forward cones of g_{eff}. This hypothesis underlies no–signalling in Part I.

H2 (Hadamard/GNS positivity).

The quadratic spectral action $\mathcal{L}^{(2)}$ defines a Green–hyperbolic kinetic operator K with causal propagators $\Delta^{\pm}$. The Peierls bracket yields the canonical (anti)commutator on the BRST physical quotient (ghost number 0):

$$[\hat{\Phi}(f), \hat{\Phi}(g)] = i\,\Delta(f,g)\mathbb{I}, \qquad \Delta := \Delta^{-} - \Delta^{+}.$$

Two–point functions are of Hadamard form, hence the GNS representation is positive and microcausal on (M, g_{eff}).

H3 (Complete positivity of reduced dynamics).

In the Markov window, reduced dynamics are generated by a GKSL Liouvillian $\mathcal{L}$ with Kossakowski matrix $C \succeq 0$:

$$\dot{\rho} = -i[H,\rho] + \sum_{\alpha,\beta} C_{\alpha\beta}\Big(L_\alpha \rho L_\beta^\dagger - \tfrac{1}{2}\{L_\beta^\dagger L_\alpha, \rho\}\Big).$$

Outside Markovianity (colored noise), there exists a *time–local CP embedding* $\dot{\rho}_t = \mathcal{L}_t[\rho_t]$ (e.g. via pseudomodes/reaction coordinates) such that the two–parameter maps $\Lambda_{t,s}$ are CP–divisible on the declared DoV.

H4 (Mixing gap on pointer sectors).

Calibrated pointer sectors are blocks selected by a POVM $\{\Pi_i\}$. Restricted to the observable algebra generated by these blocks, the Liouvillian has a spectral gap $\Delta_{\mathcal{L}} > 0$ for the nonzero part of its spectrum:

$$\big\|\Lambda_t(\rho) - \rho_\infty\big\|_1 \le C\, e^{-\Delta_{\mathcal{L}} t} \quad \text{for all initial } \rho \text{ in the sector,}$$

with ρ_∞ the sector's stationary state. The gap controls finite–time deviations in the Born–limit (Part I).

H5 (Calibrated instruments and uncertainty budgets).

Measurement devices are modeled by instruments $\mathsf{Instr} = \{\mathsf{Instr}_i\}$ with effects $E_i = \mathsf{Instr}_i^\dagger(\mathbb{I})$ (POVM). Each device carries a stated range of validity and an *uncertainty budget* for target observables A, B:

$$\varepsilon(A)^2 = \mathrm{Tr}[(M_A - A)^2\rho], \qquad \eta(B)^2 = \mathrm{Tr}[(B' - B)^2\rho],$$

together with certified error–disturbance bounds (Ozawa/Branciard–type) that are honored within the same DoV. Calibrations fix the mapping between device knobs and the theoretical parameters entering Instr.

1.2.1 DoV and Boundary Triggers

We declare a *Domain of Validity* (DoV) for each result and figure. The DoVcollects the numerical ranges and regime assumptions under which hypotheses H0–H5 and all gates Section 1.1.3 hold.

Definition 1.11 (Domain of Validity (DoV))**.** A DoV is specified by scale windows and bounds

$$\mathrm{DoV} = \Big(L_{\mathrm{loc}}^{\max},\ \Delta_{\mathcal{L}}^{\min},\ \text{noise envelopes } S(\omega),\ \text{device ranges},\ \text{background gradients}\Big),$$

such that: (i) kernels respect locality on (M, g_{eff}), i.e. $\operatorname{supp} K \subset \{(x,y) : d_{g_{\mathrm{eff}}}(x,y) \leq L_{\mathrm{loc}}^{\max}\}$; (ii) the Liouvillian on calibrated pointer sectors is CP (or CP–divisible for time–local embeddings) and admits a mixing gap $\Delta_{\mathcal{L}} \geq \Delta_{\mathcal{L}}^{\min} > 0$; (iii) stability/hyperbolicity and entropy production inequalities of Section 1.1.3 hold throughout the window.

Scales & approximations.

All asymptotic statements (Markovian windows, weak–coupling, slowly varying backgrounds) are made *within* the declared DoV. Whenever an approximation is used, we record the small parameter (e.g. $\epsilon = \|C\|/\Omega$ for a separation of time–scales) and the order kept in the expansion.

Boundary triggers (end–of–validity).

We declare an immediate end–of–validity if *any* of the following is detected for parameters inside a run:

1. **Gap closure:** $\Delta_{\mathcal{L}} \downarrow 0$ on a pointer sector (loss of exponential mixing).

2. **Loss of CP:** a negative eigenvalue of the Kossakowski matrix C, or failure of CP–divisibility for the dynamical maps $\Lambda_{t,s}$.

3. **Locality violation:** nonnegligible weight outside the g_{eff}-cones, i.e. $d_{g_{\mathrm{eff}}}(x,y) > L_{\mathrm{loc}}^{\max}$ in a kernel entering the dynamics.

4. **Device overreach:** a calibration parameter leaves its certified range, or a stated error–disturbance bound is violated within uncertainties.

This subsection replaces the terse placeholder "Scales/regimes/approximations; gap closure, loss of CP, or locality violation ⇒ end-of-validity" with a formal definition and explicit triggers, ready to be cited across Parts I–V.

1.3 Conventions and Notation

Metric, indices, units.

Spacetime carries the effective metric g_{eff} with signature $(-+++)$. Greek indices $\mu, \nu, \ldots$ denote spacetime components; lowercase Latin $a, b, \ldots$ field components; uppercase Latin $A, B, \ldots$ fiber coordinates; and adjoint indices $I, J, \ldots$ for $\mathrm{Iso}(F, h)$. We keep $\hbar$ and c explicit when clarity benefits; otherwise $\hbar = c = 1$. Raising/lowering uses g_{eff} and h.

Connections and curvatures.

∇_μ is the Levi–Civita connection of g_{eff}. An Ehresmann derivative D_μ on $E \to M$ acts on fiber fields y^A by

$$D_\mu y^A = \partial_\mu y^A + \Gamma^A{}_B(y)\, \partial_\mu y^B - \xi^A_I(y)\, A^I_\mu, \qquad \mathcal{F}^I_{\mu\nu} = \partial_\mu A^I_\nu - \partial_\nu A^I_\mu + f^I{}_{JK} A^J_\mu A^K_\nu.$$

We write the total covariant derivative on sections as $\mathbb{D}_\mu = \nabla_\mu \oplus D_\mu$, with $[\mathbb{D}_\mu, \mathbb{D}_\nu] = \mathsf{R}_{\mu\nu}(g_{\text{eff}}) \oplus \mathcal{F}_{\mu\nu}$.

Distributions and supports.

$\Delta^\pm$ are the advanced/retarded propagators of the Green–hyperbolic kinetic operator; $\Delta = \Delta^- - \Delta^+$. $J^\pm_{g_{\text{eff}}}(x)$ denote the causal future/past of x. Locality is encoded by kernel support $\operatorname{supp} K \subset \{(x, y) : d_{g_{\text{eff}}}(x, y) \le L_{\text{loc}}\}$ (see H1).

Algebra and norms.

Tr is the trace, Spec the spectrum, and $\|\cdot\|_1$ the trace norm. Projectors are Π_i with $\sum_i \Pi_i = \mathbb{I}$. Channels Λ_t are CPTP unless stated otherwise; the Liouvillian $\mathcal{L}$ generates $\Lambda_t = e^{t\mathcal{L}}$ in Markov windows.

1.4 Acronyms and Shorthands

MSP: Master Spectral Principle; DoV: Domain of Validity; GKSL: Gorini–Kossakowski–Sudarshan–L
CP/CPTP: (completely) positive / trace-preserving; POVM: positive-operator valued measure; GNS: Gelfand–Naimark–Segal; DFS/NSS: decoherence-free / noiseless subsystems.

1.5 Core Symbol Inventory

Symbol	Meaning	Typical units
g_{eff}	effective spacetime metric	–
L_{loc}	locality radius (H1)	length
$\Delta_{\mathcal{L}}$	Liouvillian mixing gap (H4)	rate
κ_G	geometric gauge factor	dimensionless
$\mathcal{N}_u, \mathcal{N}_d, \mathcal{N}_\ell$	family-count factors	integers
$\mathcal{F}_{\mu\nu}$	gauge curvature of A_μ	field strength
Π_i	pointer projectors (POVM effects)	–
ρ_∞	stationary sector state	–
μ_0	reference RG scale	energy

1.6 Proof Obligations and Artifact Policy

Canonical requirements.

We follow the *acceptance gates* in Section 1.1.3: **A1** formal Lemma/Theorem statements with explicit hypotheses and DoV; **A2** reproducible artifacts (cards/seeds/HDF5, notebooks, CI tests); **A3** compliance with Section 1.1.3. All theorems, figures, and datasets in this volume are shipped with those artifacts. (See also **??**.)

1.7 Reference Scales and Calibration Policy

Reference scale.

Coupling relations are reported at a reference scale μ_0 and then RG–transported to M_Z. Uncertainties are split as $\Delta^2_{\text{tot}} = \Delta^2_{\text{geom}} + \Delta^2_{\text{RG}}$.

Calibration and error budgets.

Devices are modeled by instruments Instr with effects E_i (POVM). Each device declares a validity range and error–disturbance budgets honored within the same DoV; knob–to–parameter maps are recorded for Zeno/filter protocols used later.

1.8 Crosswalk: Vols. I–IV to MSP Axioms

- **Vol. I** → MSP kinematics and Gate Section 1.1.3: structural fields $(\rho, \eta, \phi, T, \chi^\mu)$, fiber geometry (F, h), emergent g_{eff}, Ehresmann transport.

- **Vol. II** → Ward/BRST (Milestone M4): anomalies, background-field method, Slavnov–Taylor identities.
- **Vol. III** → calibration/positivity (Gate Section 1.1.3): emergent electromagnetism, anisotropy constraints, instrument anchoring.
- **Vol. IV** → precision/CP/entropy (Gates Section 1.1.3): spectral calculus on bundles, precision-matching maps.

1.9 DoV and Boundary Triggers (Pointer)

For the *canonical* definition of the DoV and the explicit boundary triggers (gap closure, CP loss, locality violation, device overreach), see Section 1.2.1. That DoV is assumed across Part II and cited wherever finite-time bounds and approximations are used.

Part II

Damascus Interpretation: Compact Proofs, Bounds, and Closure

Chapter 2

Hilbert Space from $\mathcal{L}^{(2)}$ and Pointer Sectors

Prelude. We construct the quantum kinematics associated with the quadratic spectral action $\mathcal{L}^{(2)}$ on the emergent background (M, g_{eff}). The route is operational: (i) a pre-symplectic current extracted from $\mathcal{L}^{(2)}$; (ii) a Cauchy-surface independent symplectic form under hyperbolicity (G); (iii) reduction by constraints/gauge (BRST ghost #0); (iv) the Peierls bracket; (v) CCR/CAR quantization and a Hadamard GNS representation (S). The hypotheses used are H0–H3 on the declared DoV.

2.1 Pre-symplectic Reduction and GNS

Setting.

Let Φ^a be a multiplet of (bosonic/fermionic) fields and $\mathcal{L}^{(2)}$ a quadratic Lagrangian density on (M, g_{eff}):

$$\mathcal{L}^{(2)}[\Phi] = \tfrac{1}{2}\,\nabla_\mu\Phi^a\, K^{\mu\nu}{}_{ab}\,\nabla_\nu\Phi^b + \tfrac{1}{2}\,\Phi^a M_{ab}\Phi^b + \Phi^a J_a, \qquad K^{\mu\nu}{}_{ab} = K^{\nu\mu}{}_{ba}, \quad M_{ab} = M_{ba}. \tag{2.1}$$

Here K is Green–hyperbolic on (M, g_{eff}) (H2,G); the Euler–Lagrange equation reads $P\Phi := -(\nabla_\mu K^{\mu\nu}\nabla_\nu)\Phi + M\Phi + J = 0$. Gauge symmetries (if present) act linearly $\delta_\epsilon\Phi = R\epsilon$. We denote by $\mathrm{Sol} \subset \Gamma(M)$ the space of smooth solutions with compactly supported Cauchy data.

Pre-symplectic current and form.

Varying $\mathcal{L}^{(2)}$ yields $\delta\mathcal{L}^{(2)} = E(\Phi)\cdot\delta\Phi + \nabla_\mu\theta^\mu(\Phi;\delta\Phi)$, with Euler form $E(\Phi) = P\Phi$. For linear theories the standard current $\omega^\mu(\delta_1\Phi, \delta_2\Phi) := \delta_1\theta^\mu(\Phi;\delta_2\Phi) - \delta_2\theta^\mu(\Phi;\delta_1\Phi)$ is field-independent and conserved on-shell: $\nabla_\mu\omega^\mu = 0$ whenever $P\Phi = 0$. The pre-symplectic

form on a Cauchy surface $\Sigma \subset M$ with future unit normal n_μ is

$$\Omega_\Sigma(\delta_1\Phi, \delta_2\Phi) = \int_\Sigma n_\mu\, \omega^\mu(\delta_1\Phi, \delta_2\Phi)\, \mathrm{d}\Sigma. \tag{2.2}$$

Theorem 2.1 (Cauchy-surface independence & Green propagators)**.** *Assume H1 (locality) and H2 (Green-hyperbolicity) on (M, g_{eff}). Then Ω_Σ is independent of Σ on the on-shell tangent space $T_\Phi\mathrm{Sol}$. Moreover, there exist unique advanced/retarded Green operators $\Delta^\pm : \Gamma_0(M) \to \Gamma(M)$ (compactly supported sources) for P, with causal propagator $\Delta := \Delta^- - \Delta^+$ obeying* $\mathrm{supp}(\Delta f) \subset J^+_{g_{\mathrm{eff}}}(\mathrm{supp}\, f) \cup J^-_{g_{\mathrm{eff}}}(\mathrm{supp}\, f)$.

Proof sketch. Green–hyperbolicity (H2,G) gives $\Delta^\pm$ with the stated support properties. Conservation of ω^μ on-shell and Stokes' theorem imply $\Omega_{\Sigma_1} - \Omega_{\Sigma_2} = \int_{\partial\mathcal{V}}\omega = 0$ for any spacetime slab $\mathcal{V}$ bounded by $\Sigma_{1,2}$. □

Reduction by constraints/gauge (BRST).

Let $\mathcal{N} := \{\delta\Phi = R\epsilon\}$ be the gauge directions and $\mathcal{C} := \overline{\mathrm{Im}\, P}$ the constraint directions. Both lie in $\ker \Omega_\Sigma$. The *kinematic phase space* is the quotient

$$V := \mathrm{Sol}\big/\big(\mathcal{N} \oplus \mathcal{C}\big), \quad \text{equivalently} \quad V \simeq H^0_{\mathrm{BRST}} \text{ (ghost number 0)}. \tag{2.3}$$

Ω_Σ descends to a nondegenerate symplectic form Ω on V (S: no ghosts/tachyons within the DoV).

Peierls bracket and the CCR/CAR algebra.

For a linear observable $F_f(\Phi) := \int_M f_a\, \Phi^a\, \mathrm{d}^4x$ with compactly supported test function f, the *Peierls bracket* is

$$\{F_f, F_g\}_{\mathrm{Peierls}} = \int_{M\times M} f_a(x)\, \Delta^{ab}(x,y)\, g_b(y)\, \mathrm{d}^4x\, \mathrm{d}^4y =: \langle f, \Delta g\rangle. \tag{2.4}$$

Quantization promotes $\Phi(f) \mapsto \hat{\Phi}(f)$ and yields the CCR/CAR:

$$[\hat{\Phi}(f), \hat{\Phi}(g)] = i\,\langle f, \Delta g\rangle\, \mathbb{I} \quad \text{(bosons)}, \qquad \{\hat{\Psi}(f), \hat{\Psi}(g)\} = \langle f, \Delta g\rangle\, \mathbb{I} \quad \text{(fermions)}. \tag{2.5}$$

Theorem 2.2 (Peierls = CCR/CAR on the BRST physical quotient)**.** *On V in (2.3), the commutator (anti-commutator) of smeared quantum fields is i times the Peierls bracket $\{\cdot,\cdot\}_{\mathrm{Peierls}}$ defined by the causal propagator Δ of P. Equivalently, the Weyl/Clifford algebra generated by $\hat{\Phi}(f)$ (resp. $\hat{\Psi}(f)$) is the unique representation of (V, Ω) compatible with locality and S.*

Proof sketch. Gauge/constraint directions lie in $\ker\Omega$, hence the quotient removes degeneracies. For linear theories $\Omega(\delta_1\Phi, \delta_2\Phi) = \langle f, \Delta g\rangle$ on solution perturbations induced by sources f, g. Canonical quantization then imposes (2.5), coinciding with Peierls (2.4). □

Hadamard GNS and positivity.

A (quasi-free) state is fixed by a two-point function Λ satisfying: (i) $\Lambda - \frac{i}{2}\Delta$ is positive on test functions (S); (ii) Λ has the Hadamard wave-front set on (M, g_{eff}) (H2). The GNS construction yields a Hilbert space $\mathcal{H}$ with cyclic vector $|\Omega\rangle$ and field operators such that $\left\langle \Omega \,\middle|\, \hat{\Phi}(f)\hat{\Phi}(g)\Omega \right\rangle = \Lambda(f, g)$. Microcausality follows from the support of Δ (Theorem 2.1).

Dependencies and acceptance.

Hypotheses used: H0–H3. *Gates:* G (hyperbolicity/support), S (positivity/no ghosts). *Acceptance (A1–A3):* formal statements & proofs (Theorems 2.1,2.2); artifacts: Green operator seeds, Hadamard test scripts, and microcausality verifications.

2.1.1 Green-hyperbolicity; causal propagators

This subsection fixes the hyperbolic framework for the quadratic operator generated by $\mathcal{L}^{(2)}$ and constructs the advanced/retarded Green operators $\Delta^{\pm}$ with sharp causal support on (M, g_{eff}).

Definition 2.3 (Green–hyperbolic operator)**.** Let P be the Euler–Lagrange operator associated with $\mathcal{L}^{(2)}$, acting on compactly supported test sections $f \in \Gamma_0(M)$. We say P is *Green–hyperbolic* on (M, g_{eff}) if there exist linear maps $\Delta^{\pm} : \Gamma_0(M) \to \Gamma(M)$ such that

$$P\,\Delta^{\pm} = \mathbb{I}_{\Gamma_0(M)}, \qquad \Delta^{\pm} P = \mathbb{I}_{\Gamma_0(M)}, \qquad \operatorname{supp}(\Delta^{\pm} f) \subset J^{\pm}_{g_{\text{eff}}}(\operatorname{supp} f) \;\; \forall f,$$

where $J^{\pm}_{g_{\text{eff}}}(\cdot)$ denote the causal future/past with respect to g_{eff}. The *causal propagator* is $\Delta := \Delta^{-} - \Delta^{+}$.

Principal symbol and hyperbolicity (Gate G).

Write the principal symbol $P(x, \xi) = Z^{\mu\nu}(x)\,\xi_\mu \xi_\nu$ (for vector/tensor fields, $Z^{\mu\nu}$ acts on components). *Hyperbolicity* holds if, for every $x \in M$, there exists a covector τ with $P(x, \tau) < 0$ such that the polynomial $P(x, \tau + \lambda\zeta)$ has only real roots in λ for all ζ. Equivalently, the characteristic cones of P are contained in the forward cones of g_{eff} and the front speed satisfies $c_{\text{char}} \leq c$.

Theorem 2.4 (Existence/uniqueness of $\Delta^{\pm}$ and support)**.** *Assume H1–H2 and Gate Section 1.1.3. Then P is Green–hyperbolic and there exist unique advanced/retarded Green operators $\Delta^{\pm}$ with:*

(i) $\operatorname{supp}(\Delta^{\pm} f) \subset J^{\pm}_{g_{\text{eff}}}(\operatorname{supp} f)$ *(sharp causal support);*

(ii) $\Delta = \Delta^{-} - \Delta^{+}$ *solves* $P\,\Delta = \Delta\,P = 0$ *on* $\Gamma_0(M)$*;*

(iii) (Formal adjoint) if $P^{\dagger} = P$ *w.r.t. the natural pairing,* $\langle f, \Delta g\rangle = -\langle \Delta f, g\rangle$.

Proof sketch. Gate G ensures energy estimates and finite propagation speed for P, hence existence of $\Delta^{\pm}$ with the stated support. Uniqueness follows from causal domain-of-dependence: if two retarded solutions differ, their difference solves the homogeneous equation with compact support in a future set, thus must vanish. □

Representation formulae.

For a compactly supported source f,

$$u^{\pm} := \Delta^{\pm} f \quad \Rightarrow \quad Pu^{\pm} = f, \qquad \operatorname{supp} u^{\pm} \subset J^{\pm}_{g_{\text{eff}}}(\operatorname{supp} f).$$

If Σ is a smooth Cauchy surface and n^{μ} its future unit normal, integration by parts gives the standard causal representation of solution differences in terms of Δ and Cauchy data on Σ.

Symplectic current and surface independence.

Let ω^{μ} be the pre-symplectic current extracted from $\mathcal{L}^{(2)}$. For two solutions produced by sources f, g, one finds

$$\Omega_{\Sigma}(\delta_f\Phi, \delta_g\Phi) = \int_{\Sigma} \mathsf{n}_{\mu}\, \omega^{\mu}(\delta_f\Phi, \delta_g\Phi)\, \mathrm{d}\Sigma = \langle f, \Delta g\rangle,$$

which is independent of the Cauchy surface (cf. Theorem 2.1). This identifies the causal propagator Δ as the symplectic form kernel on the solution space.

Compatibility with BRST reduction.

Gauge directions and constraint directions lie in $\ker \Omega$; hence Δ descends to the physical (ghost #0) quotient V defined in (2.3). This will be used in the next subsection to show *Peierls = CCR/CAR.*

Dependencies and acceptance.

Hypotheses: H1 (locality radius L_{loc}), H2 (Green–hyperbolicity). *Gates:* G (hyperbolicity/finite propagation). *Acceptance (A1–A3):* statement & proof sketch above; artifacts: numerical solvers for P, verification of causal support ($\operatorname{supp} \Delta^{\pm} f \subset J^{\pm}_{g_{\text{eff}}}(\operatorname{supp} f)$), and unit tests of formal adjointness on test grids within the declared DoV.

2.1.2 Peierls = CCR/CAR (BRST kernel at ghost #0)

We show that, on the physical (BRST) phase space, the Peierls bracket generated by the quadratic operator coincides with the canonical (anti)commutation relations after quantization. Throughout we assume H0–H3 on the declared DoV and Gate G (hyperbolicity) and S (positivity/no ghosts).

Linear observables and relative dynamics.

For a compactly supported test section f and a (bosonic or fermionic) field multiplet Φ, define the linear observable $F_f(\Phi) := \int_M f_a(x)\,\Phi^a(x)\,\mathrm{d}^4x$. Let P be the Euler–Lagrange operator from $\mathcal{L}^{(2)}$ and $\Delta^\pm$ its advanced/retarded Green operators (H2). The *retarded* (resp. *advanced*) response to a source f is $\delta_f^\pm\Phi := \Delta^\pm f$. The *Peierls bracket* of F_f, F_g is defined as the difference of retarded and advanced variations:

$$\{F_f, F_g\}_{\text{Peierls}} := \delta_f^+ F_g - \delta_f^- F_g = \int_{M\times M} f_a(x)\,\Delta^{ab}(x,y)\,g_b(y)\,\mathrm{d}^4x\,\mathrm{d}^4y =: \langle f, \Delta g\rangle. \tag{2.6}$$

The bracket is bilinear, antisymmetric (bosons) or graded–symmetric (fermions), and satisfies Jacobi on the solution space.

BRST reduction and degeneracies.

Let $R\epsilon$ be a linearized gauge variation and $\operatorname{Im} P$ the constraint directions. On any Cauchy surface Σ, the pre-symplectic form Ω_Σ from (2.2) has kernel $\ker\Omega_\Sigma = \mathcal{N}\oplus\mathcal{C}$ with $\mathcal{N} := \{\delta\Phi = R\epsilon\}$ and $\mathcal{C} := \overline{\operatorname{Im} P}$. The physical phase space is the quotient

$$V := \mathrm{Sol}\big/\big(\mathcal{N}\oplus\mathcal{C}\big) \;\simeq\; H^0_{\text{BRST}} \quad \text{(ghost number 0)}, \tag{2.7}$$

on which Ω is nondegenerate (S). For any h and ϵ, $\Delta(Ph) = 0$ and $\Delta(R\epsilon) = 0$; hence the integrand $\langle f, \Delta g\rangle$ is constant on equivalence classes modulo $\mathcal{N}\oplus\mathcal{C}$ and the Peierls bracket descends to V.

Theorem 2.5 (Peierls = CCR/CAR on the BRST physical quotient)**.** *Let H0–H3 and Gates G,S hold, and let V be as in (2.7). Then:*

(i) (Bosons) The Weyl algebra $\mathcal{W}(V,\Omega)$ generated by symbols $W([f])$ with relations $W([f])W([g]) = e^{-\frac{i}{2}\langle f,\Delta g\rangle}\,W([f+g])$ admits a regular representation by field operators $\hat\Phi(f)$ such that

$$[\hat\Phi(f), \hat\Phi(g)] \;=\; i\,\langle f, \Delta g\rangle\,\mathbb{I}. \tag{2.8}$$

(ii) (Fermions) With the (skew-)fundamental solution Δ and graded symplectic form, the CAR algebra generated by $\hat\Psi(f)$ satisfies

$$\{\hat\Psi(f), \hat\Psi(g)\} \;=\; \langle f, \Delta g\rangle\,\mathbb{I}. \tag{2.9}$$

(iii) (Microcausality) If $\operatorname{supp} f$ and $\operatorname{supp} g$ are spacelike separated in (M, g_{eff}) then $\langle f, \Delta g\rangle = 0$; hence $[\hat\Phi(f), \hat\Phi(g)] = 0$ (or $\{\hat\Psi(f), \hat\Psi(g)\} = 0$).

Proof sketch. By Theorem 2.4 (existence of $\Delta^\pm$) and Theorem 2.1, for perturbations $\delta_f^\pm\Phi = \Delta^\pm f$ and $\delta_g\Phi = \Delta g$, the pre-symplectic form satisfies $\Omega(\delta_f\Phi, \delta_g\Phi) = \langle f, \Delta g\rangle$. Gauge and constraint directions lie in $\ker\Omega$, so the pairing is well-defined on V. The Weyl

(resp. CAR) relations quantize (V, Ω), yielding (2.8) (resp. (2.9)). Support properties of Δ give microcausality. □

Weyl/CAR C^*-algebras and GNS.

Given a quasi-free Hadamard state Λ (see next subsection), the GNS triple $(\pi_\Lambda, \mathcal{H}_\Lambda, |\Omega_\Lambda\rangle)$ yields field operators with two-point function $\left\langle \Omega_\Lambda \middle| \hat{\Phi}(f)\hat{\Phi}(g)\Omega_\Lambda \right\rangle = \Lambda(f, g)$, and representation of $\mathcal{W}$ (or CAR) implementing (2.8)–(2.9). Positivity and absence of ghosts are ensured by S.

Graded Peierls bracket (bosons/fermions).

Write $|F_f| = 0$ for bosonic and 1 for fermionic functionals. The graded Peierls bracket reads

$$\{F_f, F_g\}_{\mathrm{gr}} = \langle f, \Delta g \rangle\,, \qquad \{F_f, F_g\}_{\mathrm{gr}} = -(-1)^{|F_f||F_g|}\{F_g, F_f\}_{\mathrm{gr}},$$

and satisfies the graded Jacobi identity; quantization replaces $\{\cdot,\cdot\}_{\mathrm{gr}}$ with $(i)^{-1}$ times the graded commutator.

Dependencies and acceptance.

Hypotheses: H0–H3. *Gates:* G (causal support), S (positivity/no ghosts). *Acceptance (A1–A3):* formal statement & proof sketch above; artifacts: scripts verifying $\langle f, \Delta g\rangle$ from numerically built $\Delta^{\pm}$, Weyl/CAR unit tests, and microcausality checks within the DoV.

2.1.3 Hadamard parametrix on (M, g_{eff}) with fiber couplings

We construct gauge–covariant Hadamard two–point structures on the emergent background (M, g_{eff}) for fields valued in a fiber bundle $E \to M$ with fiber (F, h) and Ehresmann derivative D_μ. Throughout we assume H0–H3 on the declared DoV, Gate G (hyperbolicity), and Gate S (positivity/no ghosts).

Laplace–type form and internal transport.

After gauge fixing on the BRST complex and reduction to the physical quotient, the Euler–Lagrange operator P associated with $\mathcal{L}^{(2)}$ can be cast (locally) in Laplace–type form on sections of E:

$$P \;=\; -(\nabla^\mu \oplus D^\mu)\,(\nabla_\mu \oplus D_\mu) \;+\; \mathsf{E}, \tag{2.10}$$

where ∇ is the Levi–Civita connection of g_{eff}, D the Ehresmann derivative with connection A_μ, and E an $\mathrm{Iso}(F, h)$-equivariant endomorphism (mass/curvature terms, gauge-fixing residue). For fermions, one replaces ∇ by the spin connection and inserts the appropriate Clifford action; the Hadamard structure is obtained by applying the first–order operator to a scalar parametrix (see remark below).

Synge world function and parallel transport.

Let $\sigma(x,y)$ be Synge's world function: half the squared geodesic distance between x and y with respect to g_{eff} within a convex normal neighborhood. Define the van Vleck–Morette determinant $\Delta_{\text{VVM}}(x,y)$ and the internal parallel transporter $\mathcal{P}(x,y)$ along the unique geodesic γ from y to x:

$$\mathcal{P}(x,y) \;=\; \mathcal{P}\exp\Big(-\int_\gamma A_\mu\,\mathrm{d}x^\mu\Big), \qquad \mathcal{P}(y,y)=\mathbb{I}, \qquad D_\mu\mathcal{P}(x,y) \text{ parallel along } \gamma. \tag{2.11}$$

We use the standard $i\epsilon$-prescription $\sigma_\epsilon(x,y)=\sigma(x,y)+i\epsilon\Big(t(x)-t(y)\Big)$ for a smooth time function t.

Hadamard ansatz (bosons).

A bosonic two–point function $\Lambda(x,y)$ is of Hadamard form if, in a convex normal neighborhood and modulo a smooth kernel,

$$H_\epsilon(x,y)=\frac{1}{8\pi^2}\,\mathcal{P}(x,y)\left[\frac{U(x,y)}{\sigma_\epsilon(x,y)} \;+\; V(x,y)\,\ln\Big(\mu^2\sigma_\epsilon(x,y)\Big) \;+\; W(x,y)\right], \tag{2.12}$$

where μ is a reference scale, U,V,W are smooth bi–sections of $\mathrm{End}(E)$, and U,V are fixed recursively by transport equations along the geodesic while W encodes state dependence.

Transport/recursion equations.

Let $\sigma_{;\mu} := \nabla_\mu\sigma$ and denote by $\Box = \nabla^\mu\nabla_\mu$. The Hadamard coefficients satisfy, with covariant derivatives extended by $\nabla\oplus D$ and curvature terms of $\mathcal{F}$ absorbed in E:

$$\text{(i) Initial data:} \qquad U(x,x)=\Delta^{1/2}_{\text{VVM}}(x,x)\,\mathbb{I}, \tag{2.13}$$

$$\text{(ii) Transport for } U: \qquad 2\,\sigma^{;\mu}\,(\nabla_\mu\oplus D_\mu)U+(\Box\sigma-4)\,U=0, \tag{2.14}$$

$$\text{(iii) Recursion for } V: \qquad 2\,\sigma^{;\mu}(\nabla_\mu\oplus D_\mu)V+(\Box\sigma-2)\,V=\Big(P-\mathsf{E}\Big)U, \tag{2.15}$$

and higher $V=\sum_{n\geq 0}V_n\,\sigma^n$ follow by matching powers of σ. Gauge covariance enters through $\mathcal{P}(x,y)$ so that $H_\epsilon(x,y)$ transforms adjointly under $\mathrm{Iso}(F,h)$.

Theorem 2.6 (Gauge–covariant Hadamard parametrix)**.** *Let P be of Laplace–type (2.10) on (M,g_{eff}) and assume H1–H3. Then there exist smooth bi–coefficients U,V,W and a parallel transporter $\mathcal{P}(x,y)$ such that the distribution H_ϵ in (2.12) satisfies*

$$P_xH_\epsilon(x,y)\equiv\delta(x,y)\,\mathbb{I}\quad(\text{mod smooth}), \qquad H_\epsilon(x,y)^\dagger=H_\epsilon(y,x),$$

and is gauge–covariant: $H_\epsilon(x,y)\mapsto g(x)H_\epsilon(x,y)g(y)^{-1}$ for $g\in\mathrm{Iso}(F,h)$. Any two choices differ by a smooth, gauge–covariant kernel.

Microlocal spectrum condition (Radzikowski form).

A (quasi–free) state is *Hadamard* iff its two–point function Λ has wave–front set

$$\mathrm{WF}(\Lambda) = \Big\{(x, k_x; y, -k_y) \in T^*M^2 \setminus \{0\} : (x, k_x) \sim (y, k_y),\ k_x \in \overline{V}_x^+\Big\},$$

i.e. k_x and k_y are cotangent–parallel along a null geodesic and future–directed. In our gauge–covariant setting this condition is imposed on the $\mathrm{End}(E)$-valued kernel after stripping the parallel transporter $\mathcal{P}(x, y)$. The singular structure of H_ϵ in (2.12) realizes this wave–front set.

Positivity and state construction.

Choose a smooth W so that $\Lambda := H_\epsilon + W$ obeys $(\Lambda - \frac{i}{2}\Delta) \geq 0$ on test sections (Gate S). Then the GNS construction yields a Hilbert space $\mathcal{H}$ and field operators $\hat{\Phi}(f)$ with two–point function Λ. Different choices of W (smooth) correspond to different Hadamard states, while all ultraviolet subtractions are fixed by U, V and the scale μ.

Fermionic remark.

For a Dirac–type operator $\not{\mathbb{D}} + m$ with spin/gauge connection, one constructs a scalar parametrix H_ϵ for $(\not{\mathbb{D}} + m)(\not{\mathbb{D}} - m)$ and sets the fermionic two–point structure $S_\epsilon = (\not{\mathbb{D}} + m)H_\epsilon$, with spin–parallel transport replacing $\mathcal{P}$. The same microlocal and positivity properties hold.

Renormalization via point–splitting.

Local covariant Wick polynomials and the stress tensor are defined by *point–splitting*:

$$:\Phi^2:(x) = \lim_{y\to x}\Big(\Phi(x)\otimes\Phi(y) - H_\epsilon(x, y)\Big), \quad \langle T_{\mu\nu}\rangle_{\mathrm{ren}} = \lim_{y\to x}\mathcal{D}_{\mu\nu}\Big(\Lambda(x, y) - H_\epsilon(x, y)\Big),$$

with a state–independent bidifferential operator $\mathcal{D}_{\mu\nu}$. Dependence on μ corresponds to standard finite renormalizations.

Dependencies and acceptance.

Hypotheses: H0–H3. *Gates:* G (causal geometry), S (positivity). *Acceptance (A1–A3):* (A1) formal statements & transport/recursion equations (2.13)–(2.15); (A2) artifacts: code/notebooks that compute U_n, V_n and implement the parallel transporter $\mathcal{P}$ for test backgrounds; (A3) numerical checks of microcausality and positivity on the declared DoV.

2.2 Stabilization by large (ρ, η): Decoherence-free Sectors

We now show how large background descriptors (ρ, η) of the spectral medium stabilize *pointer sectors* into decoherence-free (or noiseless) structures, and how the Liouvillian mixing gap $\Delta_{\mathcal{L}} > 0$ controls locking rates and finite-time errors. Throughout we assume H0–H5 on the declared DoV, and Gates S (positivity) and TQ (CP consistency).

Setting.

Let $\{\Pi_i\}$ be a calibrated POVM defining pointer sectors and $\mathcal{L}$ the (possibly time-local) GKSL generator acting on density operators:

$$\dot{\rho}_t = \mathcal{L}_t[\rho_t], \qquad \mathcal{L}_t[\rho] = -i[H_t, \rho] + \sum_{\alpha,\beta} C_{\alpha\beta}(t)\Big(L_\alpha(t)\rho L_\beta^\dagger(t) - \tfrac{1}{2}\{L_\beta^\dagger(t)L_\alpha(t), \rho\}\Big).$$

The *noise algebra* is $\mathcal{A}_t = \mathrm{alg}\{L_\alpha(t), L_\alpha^\dagger(t)\}$; its commutant is $\mathcal{A}_t'$.

2.2.1 Spectral locking rates; Lyapunov exponents; gaps

Decoherence-free subspaces/subsystems (DFS/NSS).

A projector P defines a *DFS* if $\mathcal{L}_t[P\rho P] = -i[H_{\mathrm{eff}}(t), P\rho P]$ for all ρ with support in $P\mathcal{H}$. Equivalently,

$$L_\alpha(t)P = \lambda_\alpha(t)\, P, \qquad P\, H_t\, P = H_{\mathrm{eff}}(t)P, \tag{2.16}$$

so dissipation vanishes on $P\mathcal{H}$. More generally, a *noiseless subsystem* $\mathcal{H} \simeq (\mathcal{H}_{\mathrm{code}} \otimes \mathcal{H}_{\mathrm{noisy}}) \oplus \mathcal{H}_\perp$ arises if

$$L_\alpha(t)\Big|_{\mathrm{code}\otimes\mathrm{noisy}} = \mathbb{I}_{\mathrm{code}} \otimes M_\alpha(t), \qquad H_t = H_{\mathrm{code}}(t) \otimes \mathbb{I} + \mathbb{I} \otimes H_{\mathrm{noisy}}(t), \tag{2.17}$$

i.e. the noise acts trivially on the *code* subsystem.

Theorem 2.7 (Structure via noise algebra)**.** *Let $\mathcal{A}_t = \bigoplus_J \left(\mathbb{I}_{n_J} \otimes \mathbb{M}_{m_J}\right)$ be the Wedderburn decomposition of the noise algebra on $\mathcal{H}$. Then*

$$\mathcal{H} \simeq \bigoplus_J \left(\mathbb{C}^{n_J} \otimes \mathbb{C}^{m_J}\right), \qquad \mathcal{A}_t' \simeq \bigoplus_J \left(\mathbb{M}_{n_J} \otimes \mathbb{I}_{m_J}\right),$$

and each block realizes a noiseless subsystem with code dimension n_J. Any pointer projector $P \in \mathcal{A}_t'$ is invariant under $\mathcal{L}_t$ and supports an (at least) unitary evolution generated by PH_tP.

Mixing gap and spectral locking.

Let T_1 be the trace-norm tangent subspace $\{\delta\rho : \operatorname{Tr}\delta\rho = 0\}$. Define the *mixing gap* $\Delta_{\mathcal{L}}$ of $\mathcal{L}$ restricted to the orthogonal complement of steady states:

$$\Delta_{\mathcal{L}} := \inf\{-\Re\lambda : \ \lambda \in \operatorname{Spec}(\mathcal{L}|_{\mathsf{T}_1}), \ \lambda \neq 0\}.$$

The *locking rate* γ_{lock} of the pointer statistics is the best constant such that, for $p_i(t) = \operatorname{Tr}(\Pi_i \rho_t)$,

$$\Big|p_i(t) - p_i(\infty)\Big| \ \leq \ K\, e^{-\gamma_{\text{lock}} t} \qquad \forall i, \tag{2.18}$$

with K depending on the initial condition and instrument calibration.

Proposition 2.8 (Gap $\Rightarrow$ exponential locking)**.** *Assume time-homogeneous GKSL and $\Delta_{\mathcal{L}} > 0$. Then there exists $K \geq 1$ such that $\|\rho_t - \rho_\infty\|_1 \leq K\, e^{-\Delta_{\mathcal{L}} t}$, hence $\gamma_{\text{lock}} \geq \Delta_{\mathcal{L}}$ in* (2.18)*.*

Proof sketch. Dissipativity of GKSL in the Hilbert–Schmidt metric gives $\frac{\mathrm{d}}{\mathrm{d}t}\|\delta\rho_t\|_2^2 \leq -2\,\Delta_{\mathcal{L}}\,\|\delta\rho_t\|_2^2$. Norm equivalence on finite dimensions and contractivity of CPTP maps in trace-norm yield the claim. □

Lyapunov exponents (time-local case).

For a CP-divisible time-local flow $\dot{\rho}_t = \mathcal{L}_t[\rho_t]$ with uniform gap $\inf_t \Delta_{\mathcal{L}}(t) > 0$, Oseledets' theorem defines Lyapunov exponents $\{\lambda_k\}$ on T_1 satisfying $\Re\lambda_k \leq -\inf_t \Delta_{\mathcal{L}}(t)$. Consequently,

$$\|\rho_t - \rho_\infty(t)\|_1 \ \leq \ K\, \exp\Big(-\int_0^t \Delta_{\mathcal{L}}(s)\,\mathrm{d}s\Big), \tag{2.19}$$

for some K, provided the steady-state bundle $\rho_\infty(t)$ varies slowly on the DoV.

Large (ρ, η) regime.

In the spectral medium, increasing (ρ, η) suppresses off-diagonal matrix elements in the pointer basis by enlarging the dissipative part of $\mathcal{L}$ while keeping PH_tP bounded. Operationally this pushes $\mathcal{A}_t$ towards blocks with larger m_J and fixed n_J, thus stabilizing code subspaces and enlarging $\Delta_{\mathcal{L}}$.

Checklist 2.1: PO (P1): Gap conditions and rates

- Extract $\Delta_{\mathcal{L}}$ from the spectrum of the dissipator on T_1;
- Verify $L_\alpha P = \lambda_\alpha P$ (DFS) or (2.17) (NSS);
- Report γ_{lock} and compare to the lower bound $\Delta_{\mathcal{L}}$.

2.2.2 Cauchy-surface independence

We now show that, once locking takes place, pointer statistics are independent of the choice of Cauchy surface used to read out the sector.

Heisenberg picture and conserved readouts.

Let $O = \sum_i o_i \Pi_i$ be a pointer observable. The Heisenberg adjoint evolution is $\dot{O}_t = \mathcal{L}_t^*[O_t]$, where

$$\mathcal{L}_t^*[O] = i[H_t, O] + \sum_{\alpha,\beta} C_{\alpha\beta}(t)\Big(L_\beta^\dagger(t) O L_\alpha(t) - \tfrac{1}{2}\{L_\beta^\dagger(t) L_\alpha(t), O\}\Big).$$

If $O \in \mathcal{A}_t'$ and $[H_t, O] = 0$ on a DFS/NSS block then $\dot{O}_t = 0$, hence $\mathrm{Tr}(O\rho_t)$ is time-independent after transients.

Theorem 2.9 (Cauchy-surface independence of pointer statistics)**.** *Let $R \subset M$ be a globally hyperbolic slab bounded by two Cauchy surfaces Σ_1, Σ_2. Suppose the instrument and dynamics are supported within R and satisfy H1–H5 on the* DoV*. If O belongs to the commutant $\mathcal{A}_t'$ in a DFS/NSS block, then*

$$\mathrm{Tr}\Big(O\,\rho|_{\Sigma_1}\Big) \;=\; \mathrm{Tr}\Big(O\,\rho|_{\Sigma_2}\Big),$$

and, in particular, the locked probabilities $p_i = \mathrm{Tr}(\Pi_i \rho)$ are independent of the slicing Σ.

Proof sketch. Locality (H1) allows us to represent the change of $\mathrm{Tr}(O\rho)$ between Σ_1 and Σ_2 as a spacetime integral of a divergence of a GKSL current. Inside a DFS/NSS block with $O \in \mathcal{A}_t'$ and $[H_t, O] = 0$, the adjoint generator annihilates O: $\mathcal{L}_t^*[O] = 0$. Stokes' theorem then gives equality of the two surface integrals. Locking (H4) removes transients, so the statement holds asymptotically and, by contractivity, uniformly after $t \gtrsim \gamma_{\mathrm{lock}}^{-1}$. □

Operational consequence.

Any two readout protocols that differ only by a deformation of the Cauchy surface (keeping supports within the same slab R) produce identical locked statistics. This legitimizes the replacement of time-resolved detection by surface-based readouts in the DFS/NSS regime.

Checklist 2.2: PO (P2): Invariance of readouts

- Verify $O \in \mathcal{A}_t'$ and $[H_t, O] = 0$ on the stabilized block;
- Demonstrate $\mathcal{L}_t^*[O] = 0$ numerically on the DoV;
- Compare p_i read on deformed Σ surfaces within a slab R.

Artifact 2.1: AR: Locking-rate numerics and gap extraction

We provide (i) eigensolvers for the dissipator restricted to T_1 to extract $\Delta_{\mathcal{L}}$; (ii) time-domain fits of $p_i(t)$ to obtain γ_{lock} and verify the bound $\gamma_{\text{lock}} \geq \Delta_{\mathcal{L}}$; (iii) surface-deformation benchmarks confirming Theorem 2.9 within the declared DoV.

2.3 Observable Algebra, Superselection, and Coarse-Graining

We now identify the observable algebra on the Hilbert space constructed from $\mathcal{L}^{(2)}$, extract fiber Noether charges from $\mathrm{Iso}(F,h)$ symmetry, and relate them to *pointer* observables obtained by operational coarse–graining. Superselection sectors arise from two, mutually reinforcing, structures: (i) gauge/fiber symmetries (Noether/BRST); (ii) the commutant of the noise algebra generated by GKSL Lindblad operators (DFS/NSS blocks, Theorem 2.7). Throughout we assume H0–H5 on the declared DoV and Gates $\mathsf{S}, \mathsf{G}, \mathsf{TQ}$.

2.3.1 Fiber Noether charges

Symmetry and covariant currents.

Let $\mathrm{Iso}(F,h)$ act on fiber coordinates y^A via Killing vectors $\xi_I^A(y)$ (index I in the Lie algebra $\mathfrak{iso}(F,h)$), and let the fields enter S_{master} through the Ehresmann derivative D_μ. An infinitesimal transformation with parameter $\epsilon^I(x)$ acts as

$$\delta_\epsilon y^A = \epsilon^I \xi_I^A(y), \qquad \delta_\epsilon A_\mu^I = -(\nabla_\mu \epsilon^I + f^I{}_{JK} A_\mu^J \epsilon^K). \tag{2.20}$$

The variation of the Lagrangian density is a divergence, $\delta_\epsilon \mathcal{L} = \nabla_\mu(\epsilon^I K_I^\mu)$, from which the *covariant Noether current* follows:

$$J^\mu{}_I = \frac{\partial \mathcal{L}}{\partial(D_\mu y^A)}\,\xi_I^A(y) - \frac{\partial \mathcal{L}}{\partial A_\mu^I} \quad \Rightarrow \quad (D_\mu \oplus \nabla_\mu) J^\mu{}_I = 0 \quad \text{(on-shell)}. \tag{2.21}$$

The corresponding charge on a Cauchy surface Σ with unit normal n_μ is

$$Q_I(\Sigma) = \int_\Sigma n_\mu\, J^\mu{}_I\, \mathrm{d}\Sigma, \qquad \frac{\mathrm{d}}{\mathrm{d}t} Q_I(\Sigma_t) = -\int_{\partial\Sigma_t} *J_I = 0 \ \text{ if flux vanishes.} \tag{2.22}$$

Thus Q_I is independent of the slicing when boundary flux through $\partial\Sigma$ is negligible (cf. Theorem 2.9).

Algebraic structure and constraints.

At the classical level $\{Q_I, Q_J\} = f^K{}_{IJ} Q_K$. Upon quantization (§2.1.2), the commutators realize the Lie algebra:

$$[\widehat{Q}_I, \widehat{Q}_J] = i f^K{}_{IJ}\, \widehat{Q}_K, \qquad \widehat{Q}_I^\dagger = \widehat{Q}_I. \tag{2.23}$$

Gauge generators that implement constraints are BRST–exact: $\widehat{Q}_I = \{\widehat{Q}_{\mathrm{BRST}}, \widehat{\Psi}_I\}$ on the unphysical complex; physical charges are those that survive at ghost number 0 (Theorem 2.5). Their spectral projectors define *charge superselection* sectors. The physical observable algebra $\mathcal{A}_{\mathrm{phys}}$ is the commutant of the (implemented) gauge group within $\mathcal{B}(\mathcal{H})$.

Fiber Gauss law and Ward identities.

The Gauss constraint reads $(D_\mu \oplus \nabla_\mu)\, J^\mu{}_I + (\text{field eqs.}) \equiv 0$, implying Ward identities for the generating functional and, in the background-field formalism, the Slavnov–Taylor relation (Milestone M4):

$$D_\mu \frac{\delta S_{\mathrm{master}}}{\delta A^I_\mu} + \xi^A_I(y)\, \frac{\delta S_{\mathrm{master}}}{\delta y^A} \;=\; 0.$$

These identities ensure that Q_I acts by (inner) derivations on $\mathcal{A}_{\mathrm{phys}}$ and that correlation functions respect charge conservation on (M, g_{eff}).

Charge–pointer compatibility.

Let $\mathcal{A}_{\mathrm{ptr}} = \mathrm{alg}\{\Pi_i\}$ be the pointer algebra. We say a charge Q_I is *compatible* with the pointer sectors if $[Q_I, \Pi_i] = 0$ for all i; then each sector carries a representation of the charge algebra and pointer statistics are *charge superselected*. If $\|[Q_I, \Pi_i]\| = \mathcal{O}(\epsilon)$ (small misalignment), selection rules hold up to ϵ-corrections quantified below.

2.3.2 Pointer POVMs and robustness margins

Operational coarse–graining.

Let Λ_t be the reduced evolution (GKSL or CP time–local embedding), and let $\mathsf{Instr} = \{\mathsf{Instr}_i\}$ be an instrument with effects $E_i = \mathsf{Instr}_i^\dagger(\mathbb{I})$ (POVM). The *pointer statistics* at time t are $p_i(t) = \mathrm{Tr}(E_i\, \Lambda_t(\rho_0))$. We introduce a CPTP *coarse–graining channel* $\mathcal{C} : \mathcal{B}(\mathcal{H}) \to \mathcal{A}_{\mathrm{ptr}}$ by

$$\mathcal{C}(\rho) = \sum_i \mathrm{Tr}(E_i \rho)\, \Pi_i, \qquad \Pi_i \Pi_j = \delta_{ij} \Pi_i, \quad \sum_i \Pi_i = \mathbb{I}. \tag{2.24}$$

Then $\mathrm{Tr}(\Pi_i\, \mathcal{C}(\rho)) = \mathrm{Tr}(E_i \rho)$, so all operational predictions can be made inside $\mathcal{A}_{\mathrm{ptr}}$.

Algebra/POVM consistency (O1).

We require that the pointer algebra be stable under the adjoint Heisenberg flow and compatible with charge superselection:

$$\mathcal{L}_t^*[\mathcal{A}_{\text{ptr}}] \subseteq \mathcal{A}_{\text{ptr}}, \qquad [Q_I, \Pi_i] = 0 \;\; \forall I, i. \tag{2.25}$$

The first condition is satisfied if $\Pi_i \in \mathcal{A}_t'$ (commutant of the noise algebra, cf. (2.17)); the second is automatic when Π_i are functions of conserved charges or, more generally, if Π_i are block–diagonal in the charge decomposition.

Theorem 2.10 (O1: Observable algebra / POVM consistency)**.** *Assume H1–H5 and Gates S, TQ. If the noise algebra has Wedderburn decomposition $\bigoplus_J(\mathbb{I}_{n_J} \otimes \mathbb{M}_{m_J})$ and each Π_i belongs to the commutant $\bigoplus_J(\mathbb{M}_{n_J} \otimes \mathbb{I}_{m_J})$, then $\mathcal{L}_t^*[\mathcal{A}_{\text{ptr}}] \subseteq \mathcal{A}_{\text{ptr}}$. If, in addition, $[Q_I, \Pi_i] = 0$ for all I, i, then $\mathcal{A}_{\text{ptr}} \subset \mathcal{A}_{\text{phys}}$ and pointer probabilities are charge superselected.*

Robustness margins (O2).

We quantify sensitivity to (i) device miscalibration, (ii) colored–noise mismodelling, and (iii) finite–time locking. Let θ be device knobs, $E_i(\theta)$ the effects, $S(\omega)$ the noise spectrum, and Λ_t^S the corresponding channel. Define the *margin*:

$$\mathfrak{m}_i := \sup_{\substack{\|\delta\theta\| \le \delta\theta_{\max} \\ \|S - S_0\|_{\mathsf{W}} \le \delta S_{\max} \\ t \ge t_0}} \left| \operatorname{Tr}\Big(E_i(\theta + \delta\theta)\, \Lambda_t^S(\rho_0)\Big) - \operatorname{Tr}\Big(E_i(\theta)\, \Lambda_t^{S_0}(\rho_0)\Big) \right|, \tag{2.26}$$

with $t_0 \gtrsim \gamma_{\text{lock}}^{-1}$ the post–locking time and $\|\cdot\|_{\mathsf{W}}$ a weighted noise norm (below).

Proposition 2.11 (Bounds for robustness margins)**.** *Assume CP–divisible time–local dynamics and let L_{inst} be a Lipschitz constant of $E_i(\theta)$ around θ. Let $\|\cdot\|_\diamond$ be the diamond norm. Then*

$$\left|p_i' - p_i\right| \;\le\; L_{\text{inst}}\, \|\delta\theta\| + \tfrac{1}{2}\left\|\Lambda_t^S - \Lambda_t^{S_0}\right\|_\diamond + K e^{-\gamma_{\text{lock}} t}, \tag{2.27}$$

with $K \ge 1$. Moreover, for weak colored–noise perturbations modelled by a filter function $F(\omega)$,

$$\left\|\Lambda_t^S - \Lambda_t^{S_0}\right\|_\diamond \;\le\; \int_{-\infty}^{\infty} |F(\omega)|\, |S(\omega) - S_0(\omega)|\, \mathrm{d}\omega =: \|S - S_0\|_{\mathsf{W}}, \tag{2.28}$$

so that $\mathfrak{m}_i \le L_{\text{inst}} \delta\theta_{\max} + \delta S_{\max} + K e^{-\gamma_{\text{lock}} t_0}$.

Error–disturbance with colored corrections.

For noncommuting observables A, B and a calibrated instrument, Ozawa's inequality (adapted to colored embeddings) reads

$$\varepsilon(A)\, \eta(B) + \varepsilon(A)\, \Delta(B) + \Delta(A)\, \eta(B) \;\ge\; \tfrac{1}{2}\Big|\, \langle [A, B] \rangle + \delta_{\text{col}}(A, B) \Big|, \tag{2.29}$$

where $\delta_{\mathrm{col}}(A, B)$ accounts for memory effects from the CP time–local embedding (vanishes in the Markovian window). Satisfaction of (2.29) within the DoV is part of the device certification (H5) and feeds the *instrument* contribution in (2.27).

Coarse–grained superselection and centers.

Let $\mathcal{Z}(\mathcal{A}_{\mathrm{ptr}})$ be the center of the pointer algebra. If charge projectors $\{P_\alpha\}$ and pointer projectors $\{\Pi_i\}$ commute, then the common refinement $\{P_\alpha \Pi_i\}$ labels *coarse–grained superselection* blocks. Channels and instruments that preserve the commutant structure evolve states within these blocks, making selection rules operationally sharp.

Checklist 2.3: PO (O1–O2): Algebra/POVM consistency and robustness margins

- Verify (2.25) analytically (membership in the commutant) and numerically under the declared DoV.
- Calibrate L_{inst} and compute $\|S - S_0\|_{\mathsf{W}}$ via the filter $F(\omega)$ used in the colored–Zeno design.
- Report γ_{lock} and the bound $\mathfrak{m}_i \leq L_{\mathrm{inst}} \delta\theta_{\max} + \delta S_{\max} + K e^{-\gamma_{\mathrm{lock}} t_0}$.

Artifact 2.2: AR: Robustness margins and device models

Artifacts include: (i) analytical notes for currents (2.21), charges (2.22), and algebraic checks (2.25); (ii) simulators to evaluate $\|\Lambda_t^S - \Lambda_t^{S_0}\|_\diamond$ and the filter norm (2.28); (iii) calibration scripts extracting L_{inst} and margins $\mathfrak{m}_i$; (iv) unit tests of (2.29) within the DoV.

Acceptance.

A1 formal statements & proofs: Theorems 2.10 and 2.11 with hypotheses; **A2** artifacts in **??**; **A3** compliance with Gates $\mathsf{S}, \mathsf{G}, \mathsf{TQ}, \mathsf{D}$ and the global DoV.

Chapter 3

Born's Rule from Spectral Flooding

Prelude. We develop an *energy-share* picture for measurement in the spectral medium. Pointer sectors $\{\Pi_i\}$ exchange "share" through local system–meter couplings modeled by GKSL (or CP time–local) dynamics on the declared DoV. By projecting Liouville evolution onto sector populations we obtain *balance equations* with *flooding kernels* that mix shares toward a unique stationary law. A martingale formulation (via quantum trajectories / jump unravellings) provides finite-time concentration bounds that feed the Born-limit in §??.

3.1 Energy-share Dynamics: Balance Equations, Flooding Kernels, Martingales

Let ρ_t evolve under a calibrated reduced dynamics $\dot{\rho}_t = \mathcal{L}_t[\rho_t]$, with $\mathcal{L}_t$ local and CP (H1–H3; Gates G, TQ). Fix a normalized POVM/instrument $\{\Pi_i\}$ that defines pointer sectors (H5). We consider two *shares*:

$$\text{(i) Probability share:} \quad p_i(t) := \mathrm{Tr}(\Pi_i \rho_t), \qquad \sum_i p_i(t) = 1, \tag{3.1}$$

$$\text{(ii) Energy share (relative to a pointer Hamiltonian)}: \quad s_i(t) := \frac{\mathrm{Tr}(\Pi_i H_{\mathrm{ptr}}\, \rho_t)}{\mathrm{Tr}(H_{\mathrm{ptr}} \rho_t)} \in [0,1], \qquad \sum_i s_i(t) = 1 \tag{3.2}$$

Here H_{ptr} is the calibrated meter Hamiltonian associated to the instrument (its spectrum sets the readout scale; absorbed in device models).

3.1.1 Balance equations & kernels

Projection framework.

Let P be the diagonal (sector) projection superoperator and $\mathsf{Q} := \mathbb{I} - \mathsf{P}$:

$$\mathsf{P}(\rho) = \sum_i \Pi_i \rho \Pi_i, \qquad \mathsf{Q}(\rho) = \rho - \mathsf{P}(\rho). \tag{3.3}$$

Populations are $p_i(t) = \mathrm{Tr}\left(\Pi_i \, \mathsf{P}\rho_t\right)$. The standard Nakajima–Zwanzig identity gives the exact population dynamics (time-homogeneous case shown for clarity; the time-local case is analogous with ordered exponentials):

$$\frac{\mathrm{d}}{\mathrm{d}t} \mathsf{P}\rho_t = \mathsf{P}\mathcal{L}\,\mathsf{P}\rho_t + \int_0^t \mathcal{K}(t-s)\,\mathsf{P}\rho_s \,\mathrm{d}s + \mathcal{I}_t[\mathsf{Q}\rho_0], \tag{3.4}$$

$$\mathcal{K}(\tau) := \mathsf{P}\mathcal{L}\, e^{\tau\,\mathsf{Q}\mathcal{L}}\,\mathsf{Q}\mathcal{L}\,\mathsf{P}, \qquad \mathcal{I}_t := \mathsf{P}\mathcal{L}\, e^{t\,\mathsf{Q}\mathcal{L}}\,\mathsf{Q}. \tag{3.5}$$

When the off-diagonal sector $\mathsf{Q}\rho_t$ damps with a spectral gap $\Delta_{\mathcal{L}} > 0$ (H4), the inhomogeneous term $\mathcal{I}_t[\mathsf{Q}\rho_0]$ decays exponentially and (3.4) reduces, for $t \gtrsim \Delta_{\mathcal{L}}^{-1}$, to a *closed memory equation* on populations.

Classical form (probability share).

Define the population vector $p(t) = (p_i(t))_{i=1}^n$ and the matrix-valued kernel $K(\tau)$ by

$$\Big[K(\tau)\,p\Big]_i := \mathrm{Tr}\left(\Pi_i\,\mathcal{K}(\tau)\left[\sum_j p_j \frac{\Pi_j}{\mathrm{Tr}\,\Pi_j}\right]\right),$$

where the harmless normalization $\mathrm{Tr}\,\Pi_j$ turns the diagonal block into a density for bookkeeping. Then

$$\dot{p}(t) = \int_0^t K(t-s)\,p(s)\,\mathrm{d}s \;+\; r(t), \qquad r(t) := \mathrm{Tr}(\Pi\,\mathcal{I}_t[\mathsf{Q}\rho_0]) \xrightarrow[t\to\infty]{} 0. \tag{3.6}$$

The kernel satisfies the *balance identities*

$$\sum_i K_{ij}(\tau) = 0, \qquad K_{ij}(\tau) \geq 0 \;\; (i \neq j), \qquad K_{ii}(\tau) = -\sum_{j\neq i} K_{ji}(\tau), \tag{3.7}$$

which encode conservation of total probability and positivity of flows. In the Markovian window (GKSL with time-homogeneous coefficients) the convolution collapses to a *rate generator* Q:

$$\dot{p}(t) = Q\,p(t), \qquad Q_{ij} \geq 0 \;(i \neq j), \quad \sum_i Q_{ij} = 0, \tag{3.8}$$

with Q inherited from $\mathsf{P}\mathcal{L}\mathsf{P} + \mathsf{P}\mathcal{L}\mathsf{Q}(\mathsf{Q}\mathcal{L})^{-1}\mathsf{Q}\mathcal{L}\mathsf{P}$ (Drazin inverse on the stable Q-sector).

Energy-share balance.

Let $E(t) := \mathrm{Tr}(H_{\mathrm{ptr}}\rho_t) > 0$ and $s_i(t)$ as in (3.2). Then

$$\dot{s}_i(t) = \sum_{j\neq i} \Phi_{j\to i}(t) - \sum_{j\neq i} \Phi_{i\to j}(t), \qquad \Phi_{i\to j}(t) := \frac{\mathrm{Tr}\Big(H_{\mathrm{ptr}}\Pi_j\, \mathcal{L}[\Pi_i\rho_t\Pi_i]\Big)}{E(t)} - \frac{\dot{E}(t)}{E(t)}\,\frac{\mathrm{Tr}(H_{\mathrm{ptr}}\Pi_i\rho_t)}{E(t)}\,\delta_{ij}. \tag{3.9}$$

Under standard meter calibrations (no systematic pumping between pointer levels inside $\mathcal{A}_{\mathrm{ptr}}$) the last term vanishes on average and $\Phi_{i\to j} \geq 0$ for $i \neq j$, giving an energy-share kernel identical in structure to (3.7). In practice we integrate (3.9) using the same $\mathcal{K}$ with the replacement $\Pi_i \mapsto H_{\mathrm{ptr}}^{1/2}\Pi_i H_{\mathrm{ptr}}^{1/2}$.

Stationarity and uniqueness.

If K is integrable with strictly positive total mass on every off-diagonal entry and the Q-sector is mixing with gap $\Delta_{\mathcal{L}} > 0$, then the Volterra equation (3.6) defines an asymptotically Markovian semigroup with a unique stationary distribution $\pi \in \Delta^{n-1}$ satisfying

$$\int_0^\infty K(\tau)\,\mathrm{d}\tau\;\pi = 0 \quad \text{(memory–averaged generator)}.$$

In the GKSL Markovian window, π solves $Q\,\pi = 0$ with $\Re\lambda_2(Q) > 0$.

Spectral flooding.

We call the convolution mixing in (3.6) *spectral flooding*: the kernel spreads (floods) population across sectors until $p(t) \to \pi$. The locking rate obeys $\gamma_{\mathrm{lock}} \geq \min\{\Delta_{\mathcal{L}}, \Re\lambda_2\}$, where λ_2 is the spectral gap of the (memory–averaged) generator.

3.1.2 Martingales; concentration

Jump unravelling and counting processes.

In the GKSL window, choose a Kraus–jump unravelling with jump operators $\{J_\alpha\}$ and no-jump propagation U_t. Let $N_{i\to j}(t)$ be the counting process of jumps that move population from sector i to j (defined by the support of J_α on $\Pi_i\mathcal{H} \to \Pi_j\mathcal{H}$). Its *stochastic intensity* is $\lambda_{i\to j}(t) = W_{ij}(t)\,p_i(t)$, where $W_{ij}(t) \geq 0$ are the effective rates (entries of Q_t in the time-inhomogeneous case). The *compensated process*

$$M_{i\to j}(t) := N_{i\to j}(t) - \int_0^t W_{ij}(s)\,p_i(s)\,\mathrm{d}s \tag{3.10}$$

is a martingale with bounded jumps $\Delta M \in \{0, 1\}$. Likewise, for energy share we consider weighted counts with weights given by meter-level differences $\varepsilon_j - \varepsilon_i$.

Doob martingale for sector probabilities.

Define the empirical mean $\widehat{p}_i(t) := \frac{1}{t}\int_0^t p_i(s)\,\mathrm{d}s$. Using (3.8) (or the memory-averaged generator), one can write $\widehat{p}_i(t) - \pi_i$ as a linear functional of the martingales $\{M_{k\to\ell}(t)\}$

plus exponentially decaying transients (controlled by γ_{lock}). Freedman's inequality for martingales with bounded jumps gives sub-Gaussian concentration:

Theorem 3.1 (Finite-time concentration for sector means)**.** *Assume CP-divisible dynamics with uniform locking rate $\gamma_{\text{lock}} > 0$ and bounded jump intensities $W_{ij}(t) \le W_{\max}$. Then there exist constants $c_1, c_2 > 0$ (device-calibrated) such that for all $\epsilon > 0$ and $t \ge t_0 \gtrsim \gamma_{\text{lock}}^{-1}$,*

$$\mathbb{P}\Big(\big|\widehat{p}_i(t) - \pi_i\big| > \epsilon\Big) \;\le\; 2\exp\Big[-\frac{c_1\,\epsilon^2\,t}{1 + c_2\,\epsilon}\Big] \;\le\; 2\exp\Big[-\,c\,\epsilon^2\,\gamma_{\text{lock}}\,t\Big]\,, \tag{3.11}$$

for a constant $c > 0$ depending on $W_{\max}$ and the instrument's filter function. An analogous bound holds for energy shares $s_i(t)$.

Proof sketch. Write $p(t) - \pi = e^{tQ}(p(0) - \pi) + \int_0^t e^{(t-s)Q}\,\mathrm{d}M(s)$, where $\mathrm{d}M$ stacks compensated counts. Take time averages and apply Freedman's inequality to each martingale component; contractivity of e^{tQ} yields the exponential transient. The second inequality follows by comparing the predictable quadratic variation with $\gamma_{\text{lock}}\,t$. □

Identification of the Born-limit.

In §**??** we prove that, under H0–H5, the stationary law π coincides with the operational probabilities $p_i = \mathrm{Tr}(\Pi_i \rho_\infty)$, and that the measurement readout concentrates around p_i at the sub-Gaussian rate (3.11). This is the *Born-limit* for spectral flooding.

Acceptance (for this section).

A1 Balance/memory equations (3.4)–(3.6), rate generator (3.8), energy balance (3.9), and concentration bound (3.11); **A2** Artifacts: kernels $K(\tau)$ from projected Liouvillians, Monte–Carlo jump unravellings, concentration notebooks; **A3** Gates TQ (CP), S (positivity), G (locality) hold on the declared DoV.

3.2 Stationary Limit: Theorem (Born-limit) and Hypotheses

We formalize the *Born-limit* as a stationary–mixing statement for the sector populations induced by local, CP (or CP–embedded) dynamics on the declared DoV. The hypotheses below specialize H0–H5 to the measurement setting of §3.1 and guarantee existence, uniqueness, and identification of the stationary law with pointer probabilities.

3.2.1 Hypotheses (gap/mixing, calibrated channels)

BH0 (Setting and locality).

The system–meter evolution is local on (M, g_{eff}) with radius L_{loc} (H1), and the quadratic kinematics satisfy H2 (Green–hyperbolicity). The reduced dynamics $\Lambda_{t,s}$ is either a GKSL

semigroup or admits a CP time–local embedding (H3) on the declared DoV.

BH1 (Calibrated pointer sectors).

A calibrated POVM $\{E_i\}$ with projectors Π_i defines pointer sectors (§2.3); the pointer algebra $\mathcal{A}_{\text{ptr}} = \text{alg}\{\Pi_i\}$ lies in the commutant of the noise algebra, and each Π_i is charge–compatible $[Q_I, \Pi_i] = 0$ (Theorem 2.10). We write $p_i(t) = \text{Tr}(\Pi_i \rho_t)$, $\sum_i p_i(t) = 1$.

BH2 (Primitive mixing on the sector).

There exists $t_* > 0$ and a full–rank state σ on the sector such that for all density matrices ρ

$$\Lambda_{t_*}(\rho) \;\geq\; \epsilon\,\sigma \quad \text{for some } \epsilon > 0 \qquad \text{(Doeblin/positivity–improving on the sector).} \tag{3.12}$$

Equivalently, in the GKSL window the classical generator Q of (3.8) is *irreducible* and *aperiodic*.

BH3 (Mixing gap / uniform contraction).

There is a uniform Liouvillian mixing gap $\Delta_{\mathcal{L}} > 0$ on the Q-block (§2.2.1) or, in the time–local case, an integrated gap $\int_0^t \Delta_{\mathcal{L}}(s)\,\mathrm{d}s \geq \gamma_0 t$ (with $\gamma_0 > 0$) as in (2.19). Consequently,

$$\|\Lambda_{t,s}(\rho) - \rho_\infty(s)\|_1 \;\leq\; K\,e^{-\gamma_0(t-s)}.$$

BH4 (Device certification and budgets).

The instrument Instr obeys Gate TQ (CP/CPTP consistency) and the colored Ozawa/Branciard bounds (H5) within the same DoV; its calibration defines H_{ptr} and error budgets, but does not pump probability between blocks in $\mathcal{A}_{\text{ptr}}$ on average.

Assumptions **BH0–BH4** will be referenced in the theorem below.

3.2.2 Limit and identification $p_i = \text{Tr}(\Pi_i \rho_\infty)$

Theorem 3.2 (Born-limit: existence, uniqueness, and identification)**.** *Let $\{\Lambda_{t,s}\}_{t\geq s}$ be the reduced evolution satisfying **BH0–BH4** on the declared DoV, and let $\{\Pi_i\}$ be calibrated pointer projectors. Then:*

(i) ***(Unique stationary state).*** *There exists a unique faithful stationary state ρ_∞ on the pointer sector such that $\Lambda_{t,s}(\rho_\infty) = \rho_\infty$ for all $t \geq s$, and for every initial ρ_0,*

$$\|\Lambda_{t,0}(\rho_0) - \rho_\infty\|_1 \;\leq\; K\,e^{-\gamma t}, \qquad \gamma := \min\{\Delta_{\mathcal{L}}, \gamma_0\} > 0. \tag{3.13}$$

(ii) ***(Stationary populations).*** *The population vector $p(t) = (p_i(t))$ converges exponentially to a unique $\pi \in \Delta^{n-1}$, i.e. $p(t) \to \pi$ with*

$$\|p(t) - \pi\|_1 \;\leq\; K'\,e^{-\gamma t}, \tag{3.14}$$

and, in the GKSL window, π is the unique solution of $Q\,\pi = 0$ with $\Re\lambda_2(Q) > 0$.

(iii) ***(Identification).*** *The stationary populations coincide with the pointer expectations in ρ_∞:*

$$\pi_i \;=\; \mathrm{Tr}(\Pi_i\,\rho_\infty)\,. \tag{3.15}$$

(iv) ***(Finite-time deviations).*** *The empirical means $\widehat{p}_i(t) = \frac{1}{t}\int_0^t p_i(s)\,\mathrm{d}s$ obey sub-Gaussian concentration: for all $\epsilon > 0$ and $t \geq t_0 \gtrsim \gamma^{-1}$,*

$$\mathbb{P}\Big(\Big|\widehat{p}_i(t) - \pi_i\Big| > \epsilon\Big) \;\leq\; 2\exp\Big[-c\,\epsilon^2\,\gamma\,t\Big]\,, \tag{3.16}$$

with a device–calibrated constant $c > 0$ (cf. Theorem 3.1).

Proof sketch. (i) By (3.12) (Doeblin/positivity–improving) the CP flow is *primitive* on the pointer sector, hence possesses a unique faithful stationary state and converges exponentially (Perron–Frobenius type for CP maps); the rate is bounded below by γ from the Liouvillian/sector gaps (§2.2.1).

(ii) Projecting the Liouvillian with P (cf. (3.3)) yields the Volterra equation (3.6) with a memory kernel whose total mass defines a rate generator; irreducibility implies a unique stationary π and exponential approach with spectral gap $\Re\lambda_2 > 0$. In the GKSL window this is the Markov chain (3.8).

(iii) Since $\Pi_i \in \mathcal{A}_{\mathrm{ptr}} \subset \mathcal{A}'_t$ and the flow is trace–preserving, $\frac{\mathrm{d}}{\mathrm{d}t}\,\mathrm{Tr}(\Pi_i\rho_t) = \mathrm{Tr}(\Pi_i\mathcal{L}[\rho_t])$ equals the net balance of inter-sector currents, and vanishes at stationarity; thus the limiting value is exactly $\mathrm{Tr}(\Pi_i\rho_\infty)$. Because convergence holds in trace norm (3.13), and $\|\Pi_i\| \leq 1$, we obtain the identity (3.15) and the exponential estimate (3.14).

(iv) Unravelling the CP evolution into jump processes defines compensated martingales $M_{i\to j}$ ((3.10)); Freedman/Azuma inequalities with bounded jumps yield (3.16) once transients decay at rate γ (cf. Theorem 3.1). □

Remark 3.3 (Colored noise). If $\Lambda_{t,s}$ is CP–divisible (time–local embedding) with uniform Doeblin minorization on a mesh $\{t_k\}$, the same conclusions hold replacing Q by the memory–averaged generator and γ by the minimum integrated gap over mesh steps (use submultiplicativity of Dobrushin coefficients).

Acceptance.

A1 Formal statement & proof (Theorem 3.2) with BH0–BH4; **A2** Artifacts: population kernels $K(\tau)$, Markov generators Q, Monte–Carlo unravellings and concentration notebooks; **A3** Gates S (positivity), TQ (CP), and G (locality) verified on the declared DoV.

With the Born-limit established, the next section quantifies finite-time deviation bounds as explicit functions of the locking rate and device budgets (§3.3).

3.3 Deviation Bounds for Finite Locking Rates

We quantify finite-time deviations from the stationary Born–limit when the locking rate $\gamma > 0$ (cf. (3.13)–(3.14)) is finite. Two complementary viewpoints are used: (i) *deterministic contraction* of the population dynamics with rate γ (mixing error); (ii) *stochastic concentration* of empirical readouts (shot noise) via martingale inequalities. Device/systematic budgets are inserted additively using (2.27)–(2.28).

3.3.1 Sub-Gaussian finite-time bounds

Setup and notation.

Let $p(t) = (p_i(t))$ be sector populations, π the unique stationary law, and $\gamma > 0$ the locking rate from Theorem 3.2. Define time-averaged (empirical) probabilities $\widehat{p}_i(t) = \frac{1}{t}\int_0^t p_i(s)\,\mathrm{d}s$. In the GKSL window (or for a memory-averaged generator) choose a jump unravelling with counting processes $N_{i\to j}(t)$ and intensities $\lambda_{i\to j}(t) = W_{ij}(t)\,p_i(t)$, where $0 \le W_{ij}(t) \le W_{\max}$. Let the compensated martingales be $M_{i\to j}(t)$ as in (3.10).

Theorem 3.4 (Sub-Gaussian concentration for empirical sector means)**.** *Assume **BH0–BH4**. Then there exist device–calibrated constants $\kappa_i \in (0, W_{\max}]$ and $c_0 > 0$ such that for all $\epsilon > 0$ and $t \ge t_0 \gtrsim \gamma^{-1}$,*

$$\mathbb{P}\Big(\big|\widehat{p}_i(t) - \pi_i\big| > \epsilon\Big) \;\le\; 2\exp\left[-\frac{\epsilon^2\,t}{2(\kappa_i + c_0\,\epsilon)}\right] \;\le\; 2\exp\Big[-c\,\epsilon^2\,\gamma\,t\Big]\,, \tag{3.17}$$

for some $c > 0$ depending only on $\kappa_i/W_{\max}$ and the instrument's filter. The first bound is a Freedman (Bernstein–type) inequality for martingales with bounded jumps; the second uses $\kappa_i \lesssim W_{\max} \lesssim C\,\gamma$ on the declared DoV.

Proof sketch. Represent $\widehat{p}_i(t) - \pi_i$ as a time-average of a deterministic contraction term (decaying like $e^{-\gamma(t-s)}$) plus a stochastic integral in $\mathrm{d}M$. Bounded jumps $|\Delta M_{i\to j}| \le 1$ and predictable variance $\langle M\rangle_t \le \kappa_i t$ yield Freedman's inequality, producing the first exponent. Comparing κ_i with γ along the DoV (rates are proportional to dissipative gaps) gives the simplified second exponent. □

Vector and uniform bounds.

By a union bound over $i = 1, \ldots, n$ and norm equivalence on the probability simplex,

$$\mathbb{P}(\|\widehat{p}(t) - \pi\|_1 > \epsilon) \;\le\; 2n\,\exp\left[-c\,\frac{\epsilon^2}{n^2}\,\gamma\,t\right]. \tag{3.18}$$

Sharper dimension factors are available using matrix Freedman on the block martingale that stacks $\{M_{i\to j}\}$, but (3.18) suffices for acceptance gates.

Energy-share analog.

For the energy shares $s_i(t)$ of (3.2), construct weighted counts with weights $\varepsilon_j - \varepsilon_i$ (meter levels) and repeat the martingale argument; this yields (3.17) with κ_i rescaled by $\max_{i\neq j} |\varepsilon_j - \varepsilon_i|$.

Device/systematic additions.

Combine (3.17) with the robustness margin (2.27): for $t \geq t_0$,

$$\left|\widehat{p}_i(t) - \pi_i\right| \;\leq\; \underbrace{\delta_{\text{stat}}(t,\epsilon)}_{\text{prob. term}} \;+\; \underbrace{L_{\text{inst}}\delta\theta_{\max} + \delta S_{\max} + K e^{-\gamma t_0}}_{\text{systematic term}}, \tag{3.19}$$

with probability at least $1 - 2\exp[-c\,\epsilon^2\gamma t]$.

3.3.2 Error vs. γ scaling

We separate three regimes: *mixing-limited*, *shot-noise-limited*, and *systematics-limited*.

(I) Mixing-limited (deterministic contraction).

From (3.14),

$$\|p(t) - \pi\|_1 \;\leq\; K'\,e^{-\gamma t}. \tag{3.20}$$

To achieve $\|p(t) - \pi\|_1 \leq \varepsilon$ deterministically it is enough to take

$$t \;\geq\; \frac{1}{\gamma}\log\frac{K'}{\varepsilon}. \tag{3.21}$$

(II) Shot-noise-limited (statistical averaging).

From (3.17), to obtain $|\widehat{p}_i(t) - \pi_i| \leq \varepsilon$ with confidence $1-\alpha$,

$$t \;\geq\; \frac{1}{c\,\gamma}\,\frac{1}{\varepsilon^2}\,\log\frac{2}{\alpha} \quad\Rightarrow\quad \mathrm{RMSE}(\widehat{p}_i) \;=\; \mathcal{O}\!\left(\sqrt{\frac{1}{\gamma\,t}}\right). \tag{3.22}$$

Thus, for fixed wall-clock t, larger γ (stronger locking) improves concentration as $1/\sqrt{\gamma}$.

(III) Systematics-limited (device/noise budgets).

When t is large, the error floor is set by (3.19):

$$\varepsilon_\infty \;:=\; \limsup_{t\to\infty}\left|\widehat{p}_i(t) - \pi_i\right| \;\leq\; L_{\text{inst}}\delta\theta_{\max} + \delta S_{\max}. \tag{3.23}$$

Designing colored–Zeno filters to minimize the weighted norm $\|S - S_0\|_{\mathsf{W}}$ (cf. (2.28)) directly lowers this floor.

Proposition 3.5 (Tradeoff and optimal stopping)**.** *Let total error target be ε_{tot} at confidence $1-\alpha$. Choose t as the maximum of the mixing time (3.21) and the statistical*

time from (3.22)*, then require* $L_{\text{inst}}\delta\theta_{\max} + \delta S_{\max} \leq \frac{1}{2}\varepsilon_{\text{tot}}$*. This achieves* $\mathbb{P}(\|\widehat{p}(t) - \pi\|_1 \leq \varepsilon_{\text{tot}}) \geq 1 - \alpha$*.*

Spectral-gap proxies.

In the GKSL Markov window with generator Q,

$$\gamma \;\geq\; \min\{\Delta_{\mathcal{L}}, \Re\lambda_2(Q)\}, \tag{3.24}$$

where $\Delta_{\mathcal{L}}$ is the Liouvillian gap on the Q-block and $\Re\lambda_2(Q)$ the spectral gap of the classical sector chain. For nonreversible Q, $\Re\lambda_2$ can be replaced by the pseudo-spectral gap or the Dobrushin coefficient bound.

Checklist 3.1: PO (B1–B5): Existence/uniqueness/stability

- **B1** (Existence) Verify BH0–BH4 and primitivity on sector.
- **B2** (Uniqueness) Check irreducibility/aperiodicity of Q or Doeblin minorization (3.12).
- **B3** (Stability) Extract γ via (3.24); report K, K' in (3.13)–(3.14).
- **B4** (Concentration) Validate (3.17), (3.18) from martingale logs and shot-noise simulations.
- **B5** (Budgets) Quantify $L_{\text{inst}}, \delta\theta_{\max}, \delta S_{\max}$; provide the combined bound (3.19).

Artifact 3.1: AR: Monte Carlo chains and variance budgets

We provide: (i) CTMC simulators for Q and memory-kernel samplers $K(\tau)$; (ii) martingale loggers for $N_{i\to j}(t)$ and empirical variances; (iii) notebooks producing the finite-time envelopes (3.17), (3.22), and budgets (3.19); (iv) extraction of γ via spectral fits and Dobrushin/pseudo-gap estimates.

Acceptance.

A1 Formal statements & proofs: Theorem 3.4, Proposition 3.5 with hypotheses; **A2** Artifacts in **??**; **A3** Gate compliance: S (positivity), TQ (CP), G (locality) on the declared DoV.

Chapter 4

GKSL in the Spectral Medium: Locality, CP, and No-signalling

Prelude. We derive completely positive (CP) reduced dynamics for system–meter couplings that respect the *locality radius* L_{loc} on the emergent background (M, g_{eff}). Within the declared DoV, H1–H3 ensure that microscopic, spacetime–local interactions with a spectrally mixing meter generate Gorini–Kossakowski–Sudarshan–Lindblad (GKSL) semigroups (or CP time–local embeddings for colored noise). The Kossakowski matrix arises from local meter correlations and is positive by Bochner/Herglotz positivity; Lindblad operators inherit locality from the support of the coupling kernels. Operationally, instruments (POVMs) are obtained by Naimark dilations of these CP maps and calibrated to pointer sectors $\{\Pi_i\}$ introduced earlier.

4.1 Microscopic Derivation: System–Meter Coupling, CP Semigroups, Instruments

We consider a system of interest with Hilbert space $\mathcal{H}_{\mathrm{S}}$ and a meter (bath/controller) with Hilbert space $\mathcal{H}_{\mathrm{M}}$. The joint evolution is Hamiltonian and local; the reduced system dynamics on $\mathcal{H}_{\mathrm{S}}$ is CP by construction.

4.1.1 Derivation under locality radius L_{loc}

Local microscopic model.

Let $x \mapsto \mathsf{s}_\alpha(x)$ be system densities (Heisenberg fields on $\mathcal{H}_{\mathrm{S}}$) and $x \mapsto \mathsf{b}_\alpha(x)$ meter densities (Heisenberg fields on $\mathcal{H}_{\mathrm{M}}$). The interaction is localized by kernels $u_\alpha(x)$ supported inside the g_{eff}-ball of radius L_{loc} around the device worldtube $\mathcal{W}$:

$$H_{\mathrm{int}} = \lambda \sum_\alpha \int_{\Sigma_t} \mathrm{d}^3x \; u_\alpha(x)\, \mathsf{s}_\alpha(x) \otimes \mathsf{b}_\alpha(x), \qquad \operatorname{supp} u_\alpha \subset \{x : \; d_{g_{\mathrm{eff}}}(x, \mathcal{W}) \le L_{\mathrm{loc}}\}. \tag{4.1}$$

The free Hamiltonians are H_S and H_M (stationary or slowly varying on the DoV). Locality (H1) and Green–hyperbolicity (H2) ensure finite propagation.

Born–Markov–Davies limit (Markov window).

Work in the interaction picture and scale $\lambda \to 0$, $t \mapsto \tau/\lambda^2$ (Davies weak–coupling). Assume the meter is in a stationary state ρ_M with clustering (decay of correlations) on scales $\lesssim L_{\text{loc}}$ and bounded correlation functions. Define local meter correlations

$$C_{\alpha\beta}(t) := \int \mathrm{d}^3x \mathrm{d}^3y\, u_\alpha(x) u_\beta(y)\ \mathrm{Tr}\Big(\rho_M\, \mathsf{b}_\alpha(x,t)\, \mathsf{b}_\beta(y,0)\Big), \tag{4.2}$$

$$\Gamma_{\alpha\beta}(\omega) := \int_0^\infty \mathrm{d}t\, e^{i\omega t}\, C_{\alpha\beta}(t), \qquad \Gamma(\omega) \in \mathbb{C}^{r\times r}. \tag{4.3}$$

Let $A_\alpha(\omega)$ be the spectral components of $\mathsf{S}_\alpha := \int u_\alpha(x)\, \mathsf{s}_\alpha(x)\, \mathrm{d}^3x$ w.r.t. H_S, $A_\alpha(\omega) = \sum_{\varepsilon'-\varepsilon=\omega} \Pi(\varepsilon)\, \mathsf{S}_\alpha\, \Pi(\varepsilon')$. Then, in the Markov window, the reduced dynamics ρ_t of the system obeys the GKSL master equation

$$\dot\rho_t = -i[H_S + H_{\text{LS}}, \rho_t] + \sum_\omega \sum_{\alpha,\beta} \Gamma_{\alpha\beta}(\omega)\Big(A_\beta(\omega)\rho_t A_\alpha^\dagger(\omega) - \tfrac{1}{2}\{A_\alpha^\dagger(\omega)A_\beta(\omega), \rho_t\}\Big), \tag{4.4}$$

with Lamb shift $H_{\text{LS}} = \sum_\omega \sum_{\alpha,\beta} S_{\alpha\beta}(\omega)\, A_\alpha^\dagger(\omega) A_\beta(\omega)$ given by Hilbert transforms of $C_{\alpha\beta}$.

Theorem 4.1 (Locality and CP of the Davies generator)**.** *Assume H1–H3 and that $t \mapsto C_{\alpha\beta}(t)$ is positive–definite in the Bochner sense.*[1] *Then each $\Gamma(\omega)$ is positive semidefinite, so the generator (4.4) is of GKSL form (CP and trace preserving). Moreover, $A_\alpha(\omega)$ are L_{loc}-local observables (their densities are supported in the g_{eff}-neighborhood of $\mathcal{W}$), hence the dissipator is local on (M, g_{eff}).*

Proof sketch. Positivity: Bochner positivity $\Rightarrow$ nonnegative Fourier transforms $\Gamma(\omega) \succeq 0$. Locality: the support of u_α and finite propagation imply that $A_\alpha(\omega)$ are localized in the causal hull of $\operatorname{supp} u_\alpha$, which lies within the g_{eff}-cones of radius L_{loc}. □

Colored noise (time–local CP embedding).

If $C_{\alpha\beta}(t)$ decays slowly or exhibits structure (colored noise), embed the memory into pseudomodes/reaction coordinates $\{a_k\}$ with $\dot a_k = -\kappa_k a_k + \dots$ coupled locally to the system. Eliminating the pseudomodes yields a *time–local* generator $\dot\rho_t = \mathcal{L}_t[\rho_t]$ that is CP–divisible on the DoV (Gate TQ). All locality statements above extend verbatim (the couplings remain supported in the L_{loc} neighborhood of $\mathcal{W}$).

Kossakowski matrix and geometric constraints.

The Kossakowski matrix $C_{\alpha\beta} := \sum_\omega \Gamma_{\alpha\beta}(\omega)$ inherits symmetry/positivity from $C_{\alpha\beta}(t)$ and is further restricted by the fiber symmetry $\mathrm{Iso}(F, h)$ (Noether/BRST), which block–diagonalizes the dissipator in the charge decomposition (cf. §2.3.1).

[1] For all square–integrable $f_\alpha(t)$, $\sum_{\alpha,\beta}\int \mathrm{d}t\mathrm{d}s\, \overline{f_\alpha(t)}\, C_{\alpha\beta}(t-s)\, f_\beta(s) \geq 0$.

4.1.2 Instrument maps and POVMs

Operational readout from a unitary dilation.

Couple the system locally to a meter ancilla $\mathcal{A}$ initialized in $|0\rangle$, evolve with a L_{loc}-local unitary U generated by $H_S + H_{\text{int}} + H_{\mathcal{A}}$, and perform a projective measurement $\{P_i\}$ on $\mathcal{A}$. The resulting *instrument* on the system is

$$\mathsf{Instr}_i(\rho) = \mathrm{Tr}_{\mathcal{A}}\Big[(\mathbb{I}\otimes P_i)\, U(\rho\otimes|0\rangle\langle 0|)\, U^\dagger\Big] = \sum_k M_{ik}\,\rho\, M_{ik}^\dagger, \qquad E_i := \mathsf{Instr}_i^\dagger(\mathbb{I}) = \sum_k M_{ik}^\dagger M_{ik}, \tag{4.5}$$

with E_i the POVM effects and $\sum_i E_i = \mathbb{I}$ (CPTP completeness). Locality of U (from L_{loc} and Gate G) implies the Kraus operators M_{ik} are localized in the causal hull of $\mathcal{W}$.

Compatibility with GKSL evolution.

When continuous measurement or stroboscopic sampling is performed, the unobserved meter degrees (including colored pseudomodes) generate the GKSL/CP time–local evolution Λ_t between readouts, while the readouts apply instruments Instr_i. The pointer probabilities are $p_i(t) = \mathrm{Tr}(E_i\,\Lambda_t(\rho_0))$ and, in the locking regime, $p_i \to \mathrm{Tr}(\Pi_i\rho_\infty)$ where Π_i are calibrated projectors in the pointer algebra (§2.3).

Calibration and robustness.

Calibration fixes the knob–to–parameter map $\theta \mapsto (u_\alpha, \rho_M, H_{\mathcal{A}})$ and the pointer Hamiltonian H_{ptr} entering energy shares. Robustness margins against miscalibration and colored–noise mismatch obey the bounds in (2.27)–(2.28); Ozawa/Branciard error–disturbance with colored corrections (2.29) are enforced within the DoV (H5).

Theorem 4.2 (CP instruments and POVMs under locality)**.** *Let U be generated by a Hamiltonian localized as in (4.1). Then each Instr_i in (4.5) is completely positive and the family $\{\mathsf{Instr}_i\}$ is* causal *on (M, g_{eff}): instruments supported in spacelike–separated regions commute on states supported outside their past cones (microcausality from H2/H1). The associated POVM E_i is localized in the same region, and if $E_i \in \mathcal{A}_t'$ (commutant of the noise algebra) then the pointer algebra $\mathcal{A}_{\text{ptr}} = \mathrm{alg}\{\Pi_i\}$ is invariant under the Heisenberg flow (cf. Theorem 2.10).*

Proof sketch. CP: Stinespring/Naimark dilation gives Kraus form, hence CP and completeness. Locality/causality: support of U is contained in the g_{eff}-causal hull of $\operatorname{supp} u_\alpha$; microcausality of fields and Gate G imply that two dilations with spacelike–separated supports commute on observables outside their pasts. Invariance follows from membership in the commutant. □

Summary of hypotheses/gates used.

Hypotheses: H1 (locality radius L_{loc}), H2 (Green–hyperbolicity), H3 (GKSL/CP embedding), H5 (device calibration). *Gates:* G (causal support), TQ (CP/channel consistency),

S (positivity on the physical quotient). These are the inputs for the no–signalling theorem in the next section.

Acceptance.

A1 Statements/proofs: Theorems 4.1, 4.2; **A2** Artifacts: correlation samplers $C_{\alpha\beta}(t)$, Fourier $\Gamma(\omega)$, Lindblad extraction, Kraus reconstructions, locality checks; **A3** Gate compliance $(\mathsf{G}, \mathsf{TQ}, \mathsf{S})$ on the DoV.

4.2 Theorem (No-signalling): Local GKSL $\Rightarrow$ No Superluminal Signalling

We prove that *local* reduced dynamics on (M, g_{eff})—GKSL in the Markov window or CP time–local embeddings for colored noise—do not enable superluminal signalling across spacelike–separated regions within the declared DoV. The statement is operational: outcome statistics of any instrument in a region $\mathcal{O}_B$ are independent of which instrument was chosen in a spacelike–separated region $\mathcal{O}_A$, as long as $\mathcal{O}_B$ remains outside the future cone of $\mathcal{O}_A$ during the run.

4.2.1 Spacelike separation on (M, g_{eff})

Local algebras and supports.

Let $\mathcal{A}(\mathcal{O}) \subset \mathcal{B}(\mathcal{H})$ be the C^*–algebra generated by smeared fields supported in a bounded spacetime region $\mathcal{O} \subset M$. By microcausality (H2),

$$[A, B] = 0 \quad \text{whenever} \quad \operatorname{supp} A \subset \mathcal{O}_A,\ \operatorname{supp} B \subset \mathcal{O}_B,\ \mathcal{O}_A \text{ spacelike to } \mathcal{O}_B.$$

The reduced dynamics in a device worldtube $\mathcal{W}$ has adjoint generator $\mathcal{L}^*$ of GKSL form (Markov) or time–local CP form (colored) whose Lindblad/Hamiltonian densities are supported in the g_{eff}–neighborhood $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W})$ (H1).

Causal slab geometry.

Fix a slab $R = [t_0, t_1] \times \Sigma$ and two regions $\mathcal{O}_A, \mathcal{O}_B \subset R$ with

$$J^+_{g_{\text{eff}}}(\mathcal{O}_A) \cap \mathcal{O}_B = \varnothing, \qquad J^+_{g_{\text{eff}}}(\mathcal{O}_B) \cap \mathcal{O}_A = \varnothing.$$

We say *A and B are spacelike in R.* Instruments $\{\mathsf{Instr}_i^A\}$, $\{\mathsf{Instr}_j^B\}$ are implemented by Stinespring dilations localized in $\mathcal{O}_A$ and $\mathcal{O}_B$ respectively (§4.1.2).

Theorem 4.3 (No-signalling under local GKSL)**.** *Assume H1–H3 (locality, Green–hyperbolicity, CP/GKSL) on the declared DoV. Let $\Lambda_{t,s}$ be the reduced evolution generated by a local*

GKSL Liouvillian (or a CP time–local embedding; see §4.2.2). Let $\{\mathsf{Instr}_i^A\}$ *be any instrument supported in* $\mathcal{O}_A$ *and* $\{E_j^B\}$ *a POVM supported in* $\mathcal{O}_B$*, with* $\mathcal{O}_A, \mathcal{O}_B$ *spacelike in* $R = [t_0, t_1] \times \Sigma$*. Then for every initial state* ρ *supported outside the pasts of* $\mathcal{O}_{A,B}$*,*

$$\sum_i \mathrm{Tr}\Big(E_j^B\ \Lambda_{t_1,t_0}\big[\mathsf{Instr}_i^A(\rho)\big]\Big) \;=\; \mathrm{Tr}\Big(E_j^B\ \Lambda_{t_1,t_0}[\rho]\Big) \qquad \forall j. \tag{4.6}$$

Hence the marginal statistics in B *are independent of the choice of instrument in* A*; superluminal signalling is impossible.*

Proof sketch. (*Heisenberg picture*). Let $O_B \in \mathcal{A}(\mathcal{O}_B)$. Locality (H1) implies that the Heisenberg adjoint $\mathcal{L}_A^*$ generated by Hamiltonian/Lindblad densities supported in $\mathcal{O}_A$ acts trivially on O_B while $\mathcal{O}_B$ remains spacelike to $\mathcal{O}_A$: $\mathcal{L}_A^*[O_B] = 0$ and, by microcausality (H2), the Stinespring unitary U_A of Instr^A commutes with O_B. Therefore the total adjoint evolution factorizes on commuting subalgebras and $O_B(t_1) = \Lambda_{t_1,t_0}^*[O_B]$ is independent of the choice of Instr^A. Taking Schrödinger–picture traces with E_j^B yields (4.6). (*Davies/GKSL window*). In the Markov case, $\mathcal{L}^*[\cdot] = i[H,\cdot] + \sum_\alpha (L_\alpha^\dagger \cdot L_\alpha - \frac{1}{2}\{L_\alpha^\dagger L_\alpha, \cdot\})$ with all densities supported in $\mathcal{O}_A$ gives $\mathcal{L}_A^*[O_B] = 0$ since $[O_B, H_A] = [O_B, L_\alpha] = 0$. Time–ordering extends to general Λ_{t_1,t_0}^* on the slab. □

Corollary 4.4 (Causal commutativity of instruments)**.** *Under the hypotheses of Theorem 4.3, instruments localized in spacelike-separated regions commute as channels on states supported outside their pasts, i.e.* $\mathsf{Instr}^A \circ \mathsf{Instr}^B = \mathsf{Instr}^B \circ \mathsf{Instr}^A$ *on the* DoV*, and joint outcome probabilities factorize as usual.*

Finite-speed vs. zero tails.

Because H1 enforces support strictly inside the g_{eff}–cones (not merely Lieb–Robinson tails), the statement is *sharp*: there is no exponentially small leakage outside lightcones within the DoV.

4.2.2 Colored noise via CP time-local embedding

Embedding setup.

For colored noise, extend the system by local pseudomodes $\{a_k\}$ supported in $\mathcal{N}_{L_{\text{loc}}}(\mathcal{O}_A)$ with $\dot{a}_k = -\kappa_k a_k +$ (local couplings), so that the joint evolution on system⊕pseudomodes is GKSL and *local*. Eliminating pseudomodes yields a CP–divisible, time–local generator $\mathcal{L}_t$ on the system.

Theorem 4.5 (No-signalling under CP time–local embeddings)**.** *Assume the embedded GKSL on system⊕pseudomodes is local in the slab* R *and satisfies H1–H3. Then the induced time–local evolution* $\dot{\rho}_t = \mathcal{L}_t[\rho_t]$ *on the system obeys the same no–signalling identity* (4.6) *for spacelike* $\mathcal{O}_A, \mathcal{O}_B \subset R$*.*

Proof sketch. Locality and microcausality hold on the enlarged space, so Theorem 4.3 applies before eliminating pseudomodes. The partial trace over pseudomodes commutes with local instruments in spacelike regions, preserving (4.6). □

Operational form (device level).

Let $\{\mathsf{Instr}_i^A\}$ and $\{E_j^B\}$ be calibrated as in H5, with robustness budgets satisfied on the same DoV. Then for any choice of $\{\mathsf{Instr}_i^A\}$, the *unconditioned* instrument in B produces the same distribution:

$$\sum_i p_{ij} \;=\; \sum_i \mathrm{Tr}\Big(E_j^B\ \Lambda_{t_1,t_0}\big[\mathsf{Instr}_i^A(\rho)\big]\Big) \;=\; \mathrm{Tr}\Big(E_j^B\ \Lambda_{t_1,t_0}[\rho]\Big),$$

and, if both sides perform measurements, the joint law is symmetric under $A \leftrightarrow B$ when $\mathcal{O}_A, \mathcal{O}_B$ are spacelike.

Dependencies and acceptance.

Hypotheses: H1 (locality radius L_{loc}), H2 (microcausality), H3 (GKSL/CP embedding), H5 (device calibration). *Gates:* G (causal support), TQ (CP consistency), S (positivity). **A1** Formal statements & proofs: Theorems 4.3, 4.5, Cor. 4.4; **A2** Artifacts: locality checks for L_α, commutation tests of spacelike instruments, slab simulators; **A3** Gate compliance on the declared DoV.

4.3 POVMs and Spectral Error–Disturbance (Ozawa/Branciard) with Colored Corrections

We quantify the accuracy/impact of localized instruments on noncommuting observables and adapt Ozawa–Branciard bounds to *colored* (non-Markovian) noise represented by CP time–local embeddings. Throughout, instruments $\{\mathsf{Instr}_i\}$ are realized by local Stinespring dilations (§4.1.2) inside a worldtube $\mathcal{W}$ with neighborhood $\mathcal{N}_{L_{\mathrm{loc}}}(\mathcal{W})$, and hypotheses H1–H5 hold on the declared DoV. Gates G (causality/locality), TQ (CP consistency), and S (positivity) are enforced.

Setting and notation.

Let A, B be system observables (possibly unbounded but essentially self-adjoint on a common dense domain). The meter observable is M on an ancilla prepared in $|0\rangle$. A localized unitary U couples system+ancilla; the instrument Instr_i and POVM effects E_i are given by (4.5). The *unconditioned* channel is $\mathcal{M} = \sum_i \mathsf{Instr}_i$ with Heisenberg adjoint $\mathcal{M}^*$.

State-dependent error and disturbance.

With the standard operational (Naimark) definitions, for a system state ρ

$$\varepsilon_\rho(A)^2 := \mathrm{Tr}\Big[(M^{\mathrm{out}} - A^{\mathrm{in}})^2\,(\rho\otimes|0\rangle\langle 0|)\Big], \quad \eta_\rho(B)^2 := \mathrm{Tr}\Big[(B^{\mathrm{out}} - B^{\mathrm{in}})^2\,(\rho\otimes|0\rangle\langle 0|)\Big],$$

where $X^{\mathrm{in}} = X\otimes\mathbb{I}$, $M^{\mathrm{out}} = U^\dagger(\mathbb{I}\otimes M)U$, $B^{\mathrm{out}} = U^\dagger(B\otimes\mathbb{I})U$. The intrinsic spreads are $\Delta_\rho(A)^2 := \mathrm{Tr}\Big[\Big(A - \mathrm{Tr}(A\rho)\Big)^2\rho\Big]$ and similarly for B. We also use a *spectral* (distributional) error measured on the POVM E approximating the spectral measure of A:

$$\varepsilon_\rho^{\mathrm{spec}}(A) := \sup_{\mathrm{Lip}(f)\le 1}\Big|\mathrm{Tr}\Big(f(A)\rho\Big) - \sum_i f(a_i)\;\mathrm{Tr}\Big(E_i\,\rho\Big)\Big|, \tag{4.7}$$

which is the Kantorovich–Rubinstein (Wasserstein-1) distance between the ideal and realized outcome laws when A has discrete spectrum $\{a_i\}$ (extends to Borel f in the general case).

Ozawa–type inequality (Markov window).

In the GKSL (Markov) window the usual relation holds for all ρ:

$$\varepsilon_\rho(A)\,\eta_\rho(B) + \varepsilon_\rho(A)\,\Delta_\rho(B) + \Delta_\rho(A)\,\eta_\rho(B) \;\ge\; \tfrac{1}{2}\Big|\,\mathrm{Tr}\Big([A,B]\,\rho\Big)\Big|. \tag{4.8}$$

Branciard's refinement yields a tighter quadratic bound (we record it schematically for reference):

$$\varepsilon_\rho(A)^2\,\Delta_\rho(B)^2 + \eta_\rho(B)^2\,\Delta_\rho(A)^2 + 2\,\varepsilon_\rho(A)\eta_\rho(B)\sqrt{\Delta_\rho(A)^2\Delta_\rho(B)^2 - C_\rho^2} \;\ge\; C_\rho^2, \quad C_\rho := \tfrac{1}{2}\Big|\,\mathrm{Tr}([A,B]\rho)\Big|. \tag{4.9}$$

Colored corrections via CP time–local embeddings.

When memory is present, we embed the meter as in §4.2.2 and let Λ_{t_1,t_0} denote the (CP-divisible) time–local evolution during the readout window $[t_0,t_1]$. Let $\Lambda^{\mathrm{ref}}_{t_1,t_0}$ be a calibrated Markov reference (same knobs) and define the *colored deviation* on the window by the diamond distance

$$\zeta_{\mathrm{col}} := \tfrac{1}{2}\Big\|\Lambda_{t_1,t_0} - \Lambda^{\mathrm{ref}}_{t_1,t_0}\Big\|_\diamond \;\le\; \int_{-\infty}^{\infty}|F(\omega)|\,|S(\omega) - S_0(\omega)|\,\mathrm{d}\omega =: \|S - S_0\|_{\mathsf{W}},$$

where the last inequality is the filter-norm bound (2.28) (Section 2.3). Then:

Theorem 4.6 (Ozawa with colored correction)**.** *Under H1–H5 on the DoV, for any state ρ one has*

$$\varepsilon_\rho(A)\,\eta_\rho(B) + \varepsilon_\rho(A)\,\Delta_\rho(B) + \Delta_\rho(A)\,\eta_\rho(B) \;\ge\; \tfrac{1}{2}\Big|\,\mathrm{Tr}([A,B]\rho)\Big| \;-\; \delta_{\mathrm{col}}(A,B;\rho), \tag{4.10}$$

with a correction bounded by

$$0 \;\le\; \delta_{\mathrm{col}}(A,B;\rho) \;\le\; 2\,\zeta_{\mathrm{col}}\;\min\Big\{\|A\|\Delta_\rho(B),\;\|B\|\Delta_\rho(A)\Big\} \;\le\; 2\,\|S{-}S_0\|_{\mathsf{W}}\;\min\Big\{\|A\|\Delta_\rho(B),\;\|B\|\Delta_\rho(A)\Big\}. \tag{4.11}$$

An analogous subtraction term appears in Branciard's bound (4.9)*, with* C_ρ *replaced by* $C_\rho - \delta_{\mathrm{col}}(A,B;\rho)$.

Proof sketch. Compare Heisenberg evolutions under Λ and Λ^{ref} during the readout window; their difference on any observable X is bounded in operator norm by $2\zeta_{\mathrm{col}}\|X\|$ (diamond norm duality). Apply the standard Markov proof of (4.8) to the reference channel and estimate the error introduced by replacing M^{out} and B^{out} with their colored counterparts; triangle/Cauchy–Schwarz yield (4.11). The filter bound follows from (2.28). □

Spectral error vs. disturbance tradeoff.

For discretely-valued $A = \sum_i a_i P_i$ measured by POVM $E = \{E_i\}$,

$$\varepsilon_\rho^{\mathrm{spec}}(A) \;\le\; \Big(\sum_i |a_i|\Big)\,\max_i \big\|E_i{-}P_i\big\| \;\le\; \|A\|\,\max_i \big\|E_i{-}P_i\big\|, \qquad \eta_\rho(B) \;\ge\; \big\|\mathcal{M}^*(B){-}B\big\|_\rho, \tag{4.12}$$

where $\|X\|_\rho^2 := \mathrm{Tr}(X^\dagger X\,\rho)$. Combining (4.12) with (4.10) yields a distributional (spectral) tradeoff with the same colored subtraction.

Contractivity and locality checks.

CPTP contractivity gives, for all states ρ, σ,

$$\big\|\mathcal{M}(\rho) - \mathcal{M}(\sigma)\big\|_1 \;\le\; \|\rho - \sigma\|_1, \qquad \big\|\Lambda_{t_1,t_0}(\rho) - \Lambda_{t_1,t_0}(\sigma)\big\|_1 \;\le\; \|\rho - \sigma\|_1, \tag{4.13}$$

and, by locality (H1) and microcausality (H2), instruments supported in spacelike regions commute (Cor. 4.4). These furnish the *G-checks* (Gates) listed below.

Checklist 4.1: PO (G1–G4): Locality/CP checks

- **G1 (CP & completeness).** Verify Kraus form and $\sum_i E_i = \mathbb{I}$; reconstruct $\mathcal{M}$ tomographically and confirm CPTP.
- **G2 (Contractivity).** Empirically test (4.13) on a tomographic state set; bound deviations by numerical precision.
- **G3 (Locality/causality).** Confirm supports lie in $\mathcal{N}_{L_{\mathrm{loc}}}(\mathcal{W})$; demonstrate instrument commutativity on spacelike slabs (Cor. 4.4).
- **G4 (Colored correction).** Estimate $\|S - S_0\|_{\mathsf{W}}$ from measured spectra and filters; report ζ_{col} and the bound (4.11) in the DoVsheet.

Artifact 4.1: AR: Channel tomography and contractivity tests

Artifacts include: (i) gate-set tomography (GST)/process tomography routines for $\{\mathsf{Instr}_i\}$ and $\mathcal{M}$ under colored embeddings; (ii) numerical checks of (4.13); (iii) calibration of the filter $F(\omega)$ and weighted spectral norm $\|S - S_0\|_{\mathsf{W}}$; (iv) notebooks producing the colored Ozawa/Branciard envelopes (4.10)–(4.11).

Acceptance.

A1 Formal inequalities with hypotheses: (4.8), (4.9), Theorem 4.6; **A2** Artifacts in **??**; **A3** Gate compliance $(\mathsf{G}, \mathsf{TQ}, \mathsf{S})$ and documented DoV.

Chapter 5

Colored Noise, Zeno Engineering, and Thermodynamic Ledgers

Prelude. Beyond the Markov window, meter fluctuations exhibit *colored* (time–correlated) statistics. On the declared DoV, we realize such non-Markovian effects by *time–local CP embeddings*: enlarge the system with finitely many *pseudomodes*/reaction coordinates localized inside the device worldtube, let the enlarged dynamics be GKSL (hence CP and causal), and eliminate the pseudomodes to obtain a *time–local*, CP–divisible reduced flow on the system. This section gives constructive embeddings for a large class of memory kernels and operational CP diagnostics compatible with locality (H1), Green–hyperbolicity (H2), and channel consistency (H3), with device calibration (H5). These are the building blocks for Zeno filters and the thermodynamic ledgers developed later in this chapter.

5.1 Non-Markovian Kernels → Time-local Embeddings (Pseudomodes; CP Tests)

Let the (Schrödinger) reduced state satisfy a memory equation on $[0, t]$:

$$\dot{\rho}(t) = \int_0^t \mathcal{K}(t-s)\,\rho(s)\,\mathrm{d}s \quad (+\,\text{inhom. term from initial correlations that decays on the DoV}). \tag{5.1}$$

Here $\mathcal{K}(\tau)$ is a causal superoperator kernel localized in the g_{eff}–neighborhood $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W})$ (H1), and bounded on the DoV. Our aim is to replace (5.1) by a *time–local* flow $\dot{\rho}_t = \mathcal{L}_t[\rho_t]$ obtained via a CP dilation on an enlarged space.

5.1.1 Embedding constructions

Prony/Lorentzian resolvent and pseudomodes.

Many physical kernels admit a *Prony* (sum–of–exponentials) approximation:

$$\mathcal{K}(\tau) \;\approx\; \sum_{r=1}^{R} e^{-(\kappa_r - i\Omega_r)\tau}\,\mathcal{D}_r, \qquad \kappa_r > 0, \quad \Omega_r \in \mathbb{R}, \tag{5.2}$$

where each $\mathcal{D}_r$ is a *GKSL dissipator* on the system: $\mathcal{D}_r[\rho] = \sum_k \Big(L_{rk}\rho L_{rk}^\dagger - \frac{1}{2}\{L_{rk}^\dagger L_{rk}, \rho\}\Big)$. Such expansions arise from baths with *Lorentzian* (or sum–of–Lorentzians) spectral densities and from general H^∞ model reduction of correlation functions.

System⊕pseudomodes GKSL.

Introduce R bosonic pseudomodes a_r with frequencies Ω_r and dampings κ_r (to external Markov reservoirs), and couple them locally to system operators $S_r := \sum_k c_{rk} L_{rk}$ (same support as the device):

$$H_{\text{ext}} = H_S + \sum_{r=1}^{R} \Big(\Omega_r a_r^\dagger a_r + g_r\, S_r a_r^\dagger + \overline{g}_r\, S_r^\dagger a_r\Big), \tag{5.3}$$

$$\mathcal{L}_{\text{ext}}[\Xi] = -i[H_{\text{ext}}, \Xi] + \sum_{r=1}^{R} \kappa_r \Big(a_r \Xi a_r^\dagger - \tfrac{1}{2}\{a_r^\dagger a_r, \Xi\}\Big), \tag{5.4}$$

acting on Ξ (state of system⊕pseudomodes). The enlarged dynamics is GKSL, localized in $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W})$, and respects Gates $\mathsf{G}, \mathsf{TQ}, \mathsf{S}$.

Eliminating pseudomodes reproduces the kernel.

Tracing out the pseudomodes yields an exact memory equation of the form (5.1) with kernel $\mathcal{K}(\tau) = \sum_r e^{-(\kappa_r - i\Omega_r)\tau}\,\mathcal{D}_r$ when initial system–pseudomode correlations vanish.[1] Conversely, any $\mathcal{K}$ of the form (5.2) admits an embedding (5.3)–(5.4) with suitable g_r, c_{rk}.

Theorem 5.1 (Finite–rank CP embedding for Prony kernels)**.** *Let $\mathcal{K}(\tau) = \sum_{r=1}^{R} e^{-(\kappa_r - i\Omega_r)\tau}\,\mathcal{D}_r$ with $\kappa_r > 0$ and $\mathcal{D}_r$ GKSL dissipators on the system. Then there exist couplings g_r and amplitudes c_{rk} such that the enlarged GKSL generator $\mathcal{L}_{\text{ext}}$ in (5.3)–(5.4) reproduces $\mathcal{K}$ exactly upon partial trace. In particular, the reduced flow admits a* time–local *representation $\dot{\rho}_t = \mathcal{L}_t[\rho_t]$ that is CP–divisible on the* DoV.

Proof sketch. Solve the Heisenberg–Langevin equations for a_r (linear ODEs with damping κ_r and drive S_r); substitute into the system equation to obtain a convolution with kernel $e^{-(\kappa_r - i\Omega_r)\tau}$ multiplying the bilinear dissipator generated by S_r. Matching $\{S_r\}$ to $\{L_{rk}\}$ gives $\mathcal{D}_r$. CP follows from GKSL on the enlarged space; CP *divisibility* of the reduced $\mathcal{L}_t$

[1] Nonzero initial correlations generate a standard inhomogeneous term that decays as $e^{-\kappa_{\min}\tau}$ and is included in the DoVaccounting.

holds because the elimination is Gaussian and time–local (no memory in the augmented Markov model). □

General kernels via linear ODE realizations.

More general $\mathcal{K}$ can be uniformly approximated on compact time intervals by rational transfer functions (Herglotz–Nevanlinna class), yielding a finite linear realization $\dot{y} = Ay + B\,\rho,\ \dot{\rho} = Cy$ whose elimination reproduces (5.1). Passivity/positivity of $\mathcal{K}$ can be enforced by *positive real* (PR) realizations, which admit a GKSL quantization on an extended space (ancilla qudits/oscillators) with local couplings on $\mathcal{N}_{L_{\rm loc}}(\mathcal{W})$.

Time–local generator (canonical form).

The reduced flow is $\dot{\rho}_t = \mathcal{L}_t[\rho_t]$ with

$$\mathcal{L}_t[\rho] = -i[H(t), \rho] + \sum_{\ell=1}^{m(t)} \gamma_\ell(t) \Big(L_\ell(t)\rho L_\ell^\dagger(t) - \tfrac{1}{2}\{L_\ell^\dagger(t)L_\ell(t), \rho\}\Big), \tag{5.5}$$

obtained by diagonalizing the instantaneous Kossakowski matrix. CP–divisibility on the DoVholds whenever $\gamma_\ell(t) \geq 0$ for all ℓ, t; this is guaranteed by the GKSL dilation above.

Locality.

All couplings S_r and ancilla dissipators are supported in $\mathcal{N}_{L_{\rm loc}}(\mathcal{W})$. Hence the enlarged dynamics is local in the sense of §4.1.1, and the reduced $\mathcal{L}_t$ is local as well. This ensures compatibility with the no–signalling theorem (§4.2).

5.1.2 CP diagnostics

We collect operational tests that certify complete positivity and, when needed, CP–*divisibility* for the time–local flow.

(D1) Choi positivity (global CP).

For a grid $0 = t_0 < t_1 < \cdots < t_K$, compute the Choi matrices $J_k := (\Lambda_{t_k,0} \otimes \mathbb{I})(|\Phi\rangle\langle\Phi|)$ ($|\Phi\rangle$ maximally entangled). Global CP $\Leftrightarrow$ $J_k \succeq 0$ for all k. Numerical artifacts are bounded by $\epsilon_{\rm num}$.

(D2) CP–divisibility (RHP test).

For small steps δt, form intermediate maps $\Lambda_{t+\delta t,t} = \Lambda_{t+\delta t,0} \circ \Lambda_{t,0}^{-1}$ (on the image of $\Lambda_{t,0}$); compute their Choi matrices $J(t, \delta t) \succeq 0$. Equivalently, extract the canonical rates $\gamma_\ell(t)$ in (5.5); CP–divisible iff $\gamma_\ell(t) \geq 0$ for all ℓ, t. Define the RHP non–Markovianity density $\nu_{\rm RHP}(t) = \sum_\ell (-\gamma_\ell(t))_+$.

(D3) Information backflow (BLP test).

Let $D(t) = \frac{1}{2}\|\Lambda_{t,0}(\rho_1) - \Lambda_{t,0}(\rho_2)\|_1$. A positive derivative $\dot{D}(t) > 0$ witnesses non–CP–divisible evolution. The BLP measure is $\mathcal{N}_{\rm BLP} = \int_{\dot{D}>0} \dot{D}(t)\,{\rm d}t$. For our CP–embedded flows,

$\dot{D}(t) \leq 0$ on the DoV.

(D4) Contractivity and cone tests.

Verify trace–norm contractivity $\|\Lambda_{t,0}(\rho) - \Lambda_{t,0}(\sigma)\|_1 \leq \|\rho - \sigma\|_1$ on a tomographic set (within numerical tolerance), and positivity on an operator cone spanning $\mathcal{B}(\mathcal{H}_S)$ (fast screening).

(D5) Locality audit.

Confirm supports of the (instantaneous) Lindblad densities and any instrument Kraus operators lie in $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W})$; spacelike instrument pairs commute (Cor. 4.4).

Theorem 5.2 (Gate compliance for colored embeddings)**.** *If the reduced flow arises from a local GKSL dilation (5.3)–(5.4), then (i) $J_k \succeq 0$ (global CP), (ii) $\gamma_\ell(t) \geq 0$ (CP–divisible), (iii) $\dot{D}(t) \leq 0$ (contractivity), and (iv) all locality checks pass. Hence Gates $\mathsf{TQ}, \mathsf{S}, \mathsf{G}$ are satisfied on the DoV.*

Proof sketch. (i)–(iii) follow from Stinespring dilation and complete positivity of partial trace for every subinterval; (iv) follows from H1 (support inside $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W})$) and microcausality (H2). □

Approximation error control.

When $\mathcal{K}$ is approximated by a Prony series (5.2), the propagator error on $[0, T]$ is bounded by

$$\sup_{t \leq T} \left\| \Lambda_{t,0}^{\text{true}} - \Lambda_{t,0}^{\text{embed}} \right\|_\diamond \leq \int_0^T \|\mathcal{K}(\tau) - \mathcal{K}_R(\tau)\|_{\diamond\to\diamond} \, \mathrm{d}\tau,$$

hence gate compliance is preserved up to the RHS (reported in the DoVsheet).

Artifact 5.1: AR: CP diagnostics and embedding toolkit

Artifacts include: (i) Prony/PR realizations of $\mathcal{K}$ and construction of (5.3)–(5.4); (ii) Choi tomography J_k and RHP/BLP metrics; (iii) locality auditors for supports in $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W})$; (iv) error budgets for the Prony approximation and numerical tolerances.

Acceptance.

A1 Formal statements & proofs: Theorem 5.1, Theorem 5.2; **A2** Artifacts in **??**; **A3** Gate compliance ($\mathsf{G}, \mathsf{TQ}, \mathsf{S}$) and documented DoV.

5.2 Colored-Zeno Control: Filter Functions, Convex LMIs, Robustness Margins

We design *colored–Zeno* controls that suppress the effect of colored meter noise by shaping the system's frequency response with filters compatible with locality (H1) and CP

channel consistency (H3). Our design language is: (i) *filter functions* built from localized control/toggling sequences; (ii) convex (semidefinite) programs enforcing positivity/CP constraints and hardware limits; (iii) explicit *robustness margins* that connect filter shape to the weighted spectral mismatch norm of Section 2.3.

5.2.1 Design of filter functions

Toggling frame and binary modulation.

Consider a single-qubit dephasing–type coupling for clarity (generalizations to multi-axis and multi-qubit are included below). Let $y(t) \in \{\pm 1\}$ be the *toggling* function induced by π–pulses (or strong continuous bias) inside the device worldtube $\mathcal{W}$; $y(t)$ is piecewise constant on $[0, T]$, with flips at $0 = t_0 < t_1 < \cdots < t_N = T$:

$$y(t) = (-1)^k \quad \text{for } t \in [t_k, t_{k+1}), \qquad k = 0, \ldots, N-1.$$

Define the *frequency response*

$$Y(\omega) := \int_0^T y(t)\, e^{i\omega t}\, \mathrm{d}t = \sum_{k=0}^{N-1} (-1)^k \frac{e^{i\omega t_{k+1}} - e^{i\omega t_k}}{i\omega}, \qquad F(\omega) := |Y(\omega)|^2 \ \geq 0. \tag{5.6}$$

In our notation (Sec. 2.3) the colored mismatch term satisfies

$$\zeta_{\mathrm{col}} \ \leq \ \|S - S_0\|_{\mathsf{W}} := \int_{-\infty}^{\infty} |F(\omega)|\, |S(\omega) - S_0(\omega)|\, \mathrm{d}\omega, \tag{5.7}$$

so designing F to be small on bands with large uncertainty $|S - S_0|$ reduces the correction in Theorem 4.6 and the margins (3.19).

Multi-axis and matrix filters.

For general couplings $H_{\mathrm{int}}(t) = \sum_a y_a(t)\, S_a \otimes B_a$ with control components $y_a(t)$, define the matrix response

$$Y_{ab}(\omega) := \int_0^T y_a(t)\, e^{i\omega t}\, \mathrm{d}t \int_0^T \overline{y_b(s)}\, e^{-i\omega s}\, \mathrm{d}s, \qquad F_{ab}(\omega) := Y_{ab}(\omega), \tag{5.8}$$

so that the weighted mismatch becomes $\int \mathrm{Tr}\Big(|F(\omega)|\, |S(\omega) - S_0(\omega)|\Big) \mathrm{d}\omega$. All formulas below hold with $F(\omega) \mapsto F_{ab}(\omega)$ and scalar weights replaced by positive matrices.

Trigonometric (Fourier) parameterization.

For a fixed horizon T, write the nonnegative trigonometric polynomial

$$F(\omega) = c_0 + 2 \sum_{m=1}^{M} c_m \cos(m\omega\Delta), \qquad \Delta := \tfrac{T}{M}, \quad c_m \in \mathbb{R}, \tag{5.9}$$

which exactly represents F on the discrete-time grid $t_m = m\Delta$ when $y(t)$ is piecewise constant on that grid. By Fejér–Riesz,

$$F(\omega) = v(\omega)^* Q\, v(\omega), \qquad v(\omega) = [1, e^{i\omega\Delta}, \dots, e^{iM\omega\Delta}]^\top, \quad Q \succeq 0, \tag{5.10}$$

and the Fourier coefficients are affine in Q: $c_m = \sum_{k=0}^{M-m} Q_{k,k+m}$. Thus *every* admissible F corresponds to some PSD Toeplitz–like Gram matrix Q, giving a convex design variable.

Convex filter design (nominal).

On a frequency grid $\{\omega_\ell\}$ with nonnegative weights w_ℓ, minimize the weighted mismatch against a reference spectrum S_0:

$$\begin{aligned} \min_{Q\succeq 0} \quad & J(Q) := \sum_\ell w_\ell\, F(\omega_\ell)\, W(\omega_\ell), \qquad W(\omega) := |S(\omega) - S_0(\omega)|, \\ \text{s.t.} \quad & \text{hardware \& timing constraints} \iff A\bullet Q \le b. \end{aligned} \tag{5.11}$$

Typical linear constraints include: (i) DC–notch $F(0) = 0$ (remove slow drift); (ii) bandstop or bandpass masks $F(\omega) \le \overline{F}(\omega)$ on sets; (iii) pulse–count and duty–cycle constraints (linear in the c_m), and (iv) energy/slew limits that become quadratic in y but linear in Q.

Examples (closed forms).

Choosing t_k as equally spaced (CPMG) produces a comb $F(\omega) \approx \sin^2(\omega T/2)\csc^2(\omega\Delta/2)$; Uhrig (UDD) spacing $t_k = T\sin^2\!\left(\frac{\pi k}{2N+2}\right)$ yields high-order low-frequency notch $F(\omega) = \mathcal{O}(\omega^{2N+2})$ as $\omega \to 0$. Both instances arise from specific rank–factored Q.

From filter to rates (consistency with CP).

If the colored dynamics is realized by the pseudomode dilation (§5.1.1), the canonical rates in (5.5) read $\gamma_\ell(t) = \int W_\ell(\omega, t)\, F(\omega)\, S(\omega)\, \mathrm{d}\omega$ for some nonnegative kernels W_ℓ (determined by the coupling structure). Hence $\gamma_\ell(t) \ge 0$ whenever $F(\omega) \ge 0$ and $S(\omega) \ge 0$, linking filter nonnegativity directly to CP–divisibility on the DoV.

5.2.2 LMI robustness

Uncertainty sets.

Model spectral uncertainty as either an ℓ_1–*polytope* on a grid $\{\omega_\ell\}$, $\mathcal{U}_1 := \{\delta S : |\delta S(\omega_\ell)| \le \delta_\ell\}$, or an ℓ_2–*ellipsoid*, $\mathcal{U}_2 := \{\delta S : \sum_\ell \alpha_\ell^{-1} |\delta S(\omega_\ell)|^2 \le 1\}$. We design F to minimize the worst–case mismatch $\sup_{\delta S\in\mathcal{U}} \sum_\ell |F(\omega_\ell)|\, |\delta S(\omega_\ell)|$.

Robust convex programs.

Let $f_\ell := F(\omega_\ell) = v_\ell^* Q v_\ell$ with $Q \succeq 0$.

- *Polytope $\mathcal{U}_1$ (worst–case L^1):*

$$\sup_{\delta S\in\mathcal{U}_1}\sum_\ell |f_\ell|\,|\delta S_\ell| = \sum_\ell \delta_\ell\,|f_\ell| \;\le\; \sum_\ell \delta_\ell\, t_\ell \quad \text{s.t.} \quad -t_\ell \le f_\ell \le t_\ell,$$

  which is convex with variables $(Q, \{t_\ell\})$.

- *Ellipsoid $\mathcal{U}_2$ (Cauchy–Schwarz):*

$$\sup_{\delta S\in\mathcal{U}_2}\sum_\ell |f_\ell|\,|\delta S_\ell| \;\le\; \Big(\sum_\ell \alpha_\ell |f_\ell|^2\Big)^{1/2} = \big\|D^{1/2}\mathbf{f}\big\|_2,$$

  where $D = \mathrm{diag}(\alpha_\ell)$, $\mathbf{f} = [f_\ell]$. Minimize $\|D^{1/2}\mathbf{f}\|_2$ subject to $f_\ell = v_\ell^* Q v_\ell$, an SOCP/SDP.

CP/positivity LMIs (rates and Kossakowski).

Let the instantaneous Kossakowski matrix be $C(t) = \sum_\ell w_\ell(t)\, F(\omega_\ell)\, S(\omega_\ell)\, K_\ell$, where $K_\ell \succeq 0$ encode coupling geometry. Enforce

$$C(t) \;\succeq\; 0 \quad \forall t\in\mathcal{T} \quad \Longleftarrow \quad \sum_\ell w_\ell(t)\, S(\omega_\ell)\, v_\ell^* Q v_\ell\; K_\ell \;\succeq\; 0, \quad \forall t\in\mathcal{T}, \tag{5.12}$$

which is linear in Q (each t gives one LMI block). If rates $\gamma_\ell(t) = \int W_\ell F S$ are explicit linear functionals of Q, then $\gamma_\ell(t) \ge 0$ become linear inequalities.

S–procedure (robust CP).

For ellipsoidal uncertainty $\mathcal{U}_2$, require $C(t;\delta S) \succeq 0$ for all $\delta S \in \mathcal{U}_2$. A sufficient LMI certificate is

$$\exists\, \tau(t) \ge 0: \quad \begin{bmatrix} \sum_\ell w_\ell(t)\, S_0(\omega_\ell)\, f_\ell K_\ell & * \\ * & \tau(t)\, I \end{bmatrix} \succeq \sum_\ell \begin{bmatrix} 0 & \sqrt{\alpha_\ell}\, K_\ell^{1/2} f_\ell \\ \sqrt{\alpha_\ell}\, f_\ell K_\ell^{1/2} & I \end{bmatrix}, \qquad f_\ell = v_\ell^* Q v_\ell, \tag{5.13}$$

which guarantees positivity for all δS in the ellipsoid (standard S–procedure bound; tighten by gridding t).

Putting it together (master SDP).

Collect constraints into a single SDP:

$$\min_{Q\succeq 0,\; t,\; \{\tau(t)\}} \underbrace{\sum_\ell w_\ell\, |v_\ell^* Q v_\ell|\, W(\omega_\ell)}_{\text{nominal}} + \lambda \underbrace{\sup_{\delta S\in\mathcal{U}} \sum_\ell |v_\ell^* Q v_\ell|\,|\delta S_\ell|}_{\text{robust}} \tag{5.14}$$

$$\text{s.t. hardware/timing: } A\bullet Q \le b; \quad \text{CP/PSD: (5.12), (5.13)}; \quad F(0) = 0 \text{ (optionally)},$$

with tradeoff $\lambda \ge 0$. The solution yields a filter $F(\omega)$ that minimizes worst–case colored mismatch while preserving CP and implementability.

Robustness margins and stopping rules.

With the designed F, the colored correction satisfies $\zeta_{\mathrm{col}} \leq \|S - S_0\|_{\mathsf{W}}$ from (5.7), so the pointer margins in (3.19) become

$$\mathfrak{m}_i \;\leq\; L_{\mathrm{inst}}\delta\theta_{\max} + \underbrace{\int |F(\omega)|\,|S(\omega) - S_0(\omega)|\,\mathrm{d}\omega}_{\text{colored term}} + Ke^{-\gamma t_0}. \tag{5.15}$$

A practical *stopping rule*: halt averaging at the t where the statistical term (Theorem 3.4) drops below the systematic floor given by the RHS.

Checklist 5.1: Design & robustness (filters & LMIs)

- Parameterize $F(\omega) = v^* Q v$, $Q \succeq 0$ (Fejér–Riesz).
- Impose CP/PSD LMIs on rates/Kossakowski; add hardware linear constraints.
- Optimize nominal+robust objectives (5.14); export $y(t)$.
- Report $\|S - S_0\|_{\mathsf{W}}$, ζ_{col}, and margins via (5.15); verify no–signalling slicewise.

Artifact 5.2: AR: Filter synthesis and SDP notebooks

Artifacts include: (i) Gram–SDP builders for Q and time–grid mappers; (ii) libraries for CPMG/UDD/optimized sequences; (iii) CP/PSD LMI packs for rates/Kossakowski; (iv) robust solvers for (5.14); (v) margin dashboards implementing (5.15) and the stopping rule.

Acceptance.

A1 Formal constructions: filter factorization (5.10), design program (5.11), CP LMIs (5.12)–(5.13); **A2** Artifacts in **??**; **A3** Gates $\mathsf{G}, \mathsf{TQ}, \mathsf{S}$ verified on the declared DoV.

5.3 Work/Heat Ledgers: First/Second-law Balance; Landauer-type Bounds

We formulate *thermodynamic ledgers* for local GKSL and the CP time–local embeddings built in §5.1. On the declared DoV, all couplings are supported in $\mathcal{N}_{L_{\mathrm{loc}}}(\mathcal{W})$ (H1), dynamics is Green–hyperbolic (H2) and CP (H3), and instruments are calibrated (H5). The ledger records, for each time window and channel, the signed energy and entropy exchanges so that the first and second laws hold *term by term* (up to quantified colored corrections).

5.3.1 Ledger construction

System, meters, and pseudomodes.

Let the enlarged (Markov) model of §5.1.1 evolve under $\mathcal{L}_{\text{ext}}$ on system $\otimes$ pseudomodes $\{a_r\}$, each damped to a Markov reservoir $\mathcal{R}_r$ at inverse temperature β_r. The reduced system state ρ_t obeys a time–local CP generator $\mathcal{L}_t$ (canonical form (5.5)).

Internal energy, power, and heat currents.

For a (possibly driven) system Hamiltonian $H(t)$,

$$U(t) := \operatorname{Tr}\Big(H(t)\rho_t\Big), \qquad \dot{U}(t) = \underbrace{\operatorname{Tr}(\dot{H}(t)\rho_t)}_{P(t)} + \underbrace{\operatorname{Tr}\Big(H(t)\,\mathcal{L}_t[\rho_t]\Big)}_{\dot{Q}(t)}.$$

In the pseudomode dilation the dissipator decomposes into physical channels indexed by r (bath/mode label) or by Lindblad labels ℓ; write

$$\mathcal{L}_t = \sum_r \mathcal{L}_t^{(r)} \quad \Rightarrow \quad \dot{Q}(t) = \sum_r J_r(t), \qquad J_r(t) := \operatorname{Tr}\Big(H(t)\,\mathcal{L}_t^{(r)}[\rho_t]\Big).$$

Thus the *first-law ledger* on any window $[t_0, t_1]$ is

$$\Delta U \;=\; W \;+\; \sum_r Q_r, \qquad \Delta U := U(t_1) - U(t_0), \;\; W := \int_{t_0}^{t_1} P(t)\,\mathrm{d}t, \;\; Q_r := \int_{t_0}^{t_1} J_r(t)\,\mathrm{d}t. \tag{5.16}$$

Reference NESS and entropy production.

Let ρ_t^* be a *reference* instantaneous stationary state for $\mathcal{L}_t$ (e.g. the NESS selected by the pseudomode reservoirs, constant if $\mathcal{L}_t$ is time-homogeneous). Define the relative entropy $D(\rho\|\sigma) := \operatorname{Tr}\Big[\rho(\ln\rho - \ln\sigma)\Big]$. Spohn's inequality for GKSL (and CP–divisible $\mathcal{L}_t$) gives the *entropy production rate*

$$\sigma(t) := -\frac{\mathrm{d}}{\mathrm{d}t} D\Big(\rho_t\|\rho_t^*\Big) = \dot{S}(t) - \sum_r \beta_r J_r(t) \;\; \geq \;\; 0, \qquad S(t) := -\operatorname{Tr}(\rho_t \ln \rho_t), \tag{5.17}$$

with equality iff $\rho_t = \rho_t^*$ and detailed balance holds channelwise.

Ledger entries.

Fix a partition of $[t_0, t_1]$ into slices where knobs are constant. For each slice and channel r record the tuple

$$\Big(\Delta U;\, W;\, \{Q_r\}_r;\, \Delta S;\, \Sigma\Big) := \Big(\textstyle\int \operatorname{Tr}(H\,\mathrm{d}\rho) + \int \operatorname{Tr}(\dot{H}\,\rho) \;\;,\; \int \operatorname{Tr}(\dot{H}\,\rho) \;\;,\; \{\int J_r\}_r \;\;,\; \int \dot{S} \;\;,\; \int \sigma\Big),$$

which by (5.16)–(5.17) satisfies $\Delta U = W + \sum_r Q_r$ and $\Sigma = \Delta S - \sum_r \beta_r Q_r \geq 0$.

Measurement steps.

A (localized) instrument $\{\mathsf{Instr}_i\}$ with dilation U and meter Hamiltonian H_M contributes no explicit *work* if H_M and the coupling are time-independent on the slice; the energetic effect registers as heat into the meter's reservoir $r = M$ via the corresponding Lindblad part of the dilation. Hence the ledger remains valid with r ranging over physical baths and meter sinks.

Colored corrections.

When the true colored evolution Λ_{t_1,t_0} deviates from a calibrated Markov reference by ζ_{col} (diamond norm), energy and entropy balances acquire controlled additive slacks. For any observable X with $\|X\| \le \Xi$,

$$\left|\mathrm{Tr}\Big(X\,(\Lambda_{t_1,t_0} - \Lambda^{\mathrm{ref}}_{t_1,t_0})(\rho)\Big)\right| \;\le\; 2\,\Xi\,\zeta_{\mathrm{col}}, \tag{5.18}$$

hence the ledger computed with the reference channel is valid up to $\mathcal{O}(\zeta_{\mathrm{col}})$ corrections; see bounds below.

5.3.2 Bounds and consistency

Clausius and Hatano–Sasa forms.

Integrating (5.17) on $[t_0, t_1]$,

$$\Sigma \;:=\; \int_{t_0}^{t_1} \sigma(t)\,\mathrm{d}t \;=\; \Delta S \;-\; \sum_r \beta_r Q_r \;\ge\; 0. \tag{5.19}$$

If ρ_t^* varies (driving), the excess part is $\Sigma_{\mathrm{ex}} = \Delta S - \sum_r \beta_r Q_r - \int \mathrm{Tr}(\dot\rho_t^*\,\Xi_t)$, Hatano–Sasa–type, with $\Xi_t = -\ln \rho_t^*$.

Landauer for erasure and coarse readout.

For an isothermal reservoir at β and a logical map $\mathcal{E}$ (e.g. erasure or coarse-graining to pointer sectors),

$$\beta\,Q \;\ge\; \Delta S\Big(\mathcal{E}[\rho]\Big) - \Delta S(\rho) \;=:\; -\Delta H_{\mathrm{Shannon}} \;\;\Rightarrow\;\; Q \;\ge\; k_B T\,\Delta I, \tag{5.20}$$

with $\Delta I := H_{\mathrm{in}} - H_{\mathrm{out}}$ the information lost (Shannon) to the environment. Equality requires quasi-static, reversible operation in the Markov limit.

Landauer with colored correction.

Let the true colored channel on $[t_0, t_1]$ be Λ and the calibrated Markov reference be Λ^{ref}. Then

$$\Delta S_{\mathrm{true}} - \beta Q_{\mathrm{true}} \;\ge\; -\varepsilon_L, \qquad \varepsilon_L \;\le\; 2\,\zeta_{\mathrm{col}}\left(\sup_{t\in[t_0,t_1]} \|H(t)\| \;+\; \sup_{\rho} \|\ln\rho\|_\rho\right),$$

hence the Landauer inequality (5.20) holds with an additive slack $\varepsilon_L = \mathcal{O}(\zeta_{\mathrm{col}})$. When filters satisfy (5.7), ε_L is bounded by the designed weighted spectral mismatch $\|S - S_0\|_{\mathsf{W}}$.

Free-energy and work bounds.

Define the nonequilibrium free energy relative to ρ_t^*: $F_t(\rho) := \mathrm{Tr}(H(t)\rho) - \beta^{-1}S(\rho) - F_t(\rho_t^*)$ (zero at ρ_t^*). For an isothermal bath,

$$W \;\geq\; \Delta F \;+\; \beta^{-1}\,\Sigma \;\;\geq\;\; \Delta F, \qquad \Delta F := F_{t_1}(\rho_{t_1}) - F_{t_0}(\rho_{t_0}). \tag{5.21}$$

Thus the ledger yields *minimal work* for state transformations; colored corrections contribute at most $\mathcal{O}(\zeta_{\mathrm{col}})$ to ΔF via (5.18).

Rate consistency and CP.

With the canonical form (5.5), choose the quadratic form $E_\ell(t) := \mathrm{Tr}\Big(L_\ell^\dagger(t)L_\ell(t)\rho_t\Big) \geq 0$. Then each channel's heat current can be written as $J_\ell(t) = \gamma_\ell(t)\,\xi_\ell(t)$ for some bounded $\xi_\ell(t)$ depending on $H(t)$ and $L_\ell(t)$. CP–divisibility ($\gamma_\ell \geq 0$) implies that heat cannot flow in the "wrong sign" for a detailed–balance channel; checking $\gamma_\ell(t) \geq 0$ is thus a *consistency test* for the ledger.

Filter-induced cooling/heating bands.

For pseudomode r with Lorentzian spectrum centered at Ω_r and filter $F(\omega)$, the cycle-averaged heat obeys

$$Q_r \;\propto\; \int F(\omega)\,S_r(\omega)\,\Phi_r(\omega)\,\mathrm{d}\omega,$$

with $\Phi_r(\omega)$ a device geometry factor. Bandstops around Ω_r ($F \approx 0$) suppress Q_r (Zeno "cooling"); bandpasses enhance it ("heating"). The same F enters the colored correction ζ_{col} via (5.7), linking control directly to both heat and error budgets.

Ledger accuracy under model mismatch.

Let the implemented filter be $\widetilde{F}$ while design assumed F. Then the *heat ledger error* per channel satisfies

$$|Q_r - \widetilde{Q}_r| \;\leq\; \int |F(\omega) - \widetilde{F}(\omega)|\,|S_r(\omega)|\,K_r(\omega)\,\mathrm{d}\omega, \tag{5.22}$$

for a known kernel K_r determined by couplings and $H(t)$. This is the *robustness margin* reported alongside the ledger.

Checklist 5.2: PO (N1, Z1, L1): Embedding / robustness / ledger

- **N1 (Embedding)**: certify CP time–local embedding (RHP/BLP tests) and extract $\{\gamma_\ell(t)\}$; verify $\gamma_\ell(t) \geq 0$.
- **Z1 (Filters)**: record $F(\omega)$ and weighted mismatch $\|S - S_0\|_{\mathsf{W}}$; compute ζ_{col}

and bounds (5.7), (5.22).

- **L1 (Ledger)**: report $(\Delta U, W, \{Q_r\}, \Delta S, \Sigma)$ on each slice; verify (5.16) and (5.19) within the stated $\mathcal{O}(\zeta_{\mathrm{col}})$ tolerances.

Artifact 5.3: AR: Sequences/code and CP tests

Artifacts include: (i) simulators for $J_r(t)$ from $\mathcal{L}_t$ and from the pseudomode dilation; (ii) GST/process tomography for $\Lambda_{t,0}$ and Choi–based CP/RHP diagnostics; (iii) filter synthesis notebooks generating $F(\omega)$ and margins (5.22); (iv) automated ledgers that output $\Delta U, W, \{Q_r\}, \Delta S, \Sigma$ and verify the identities with colored slack computed from ζ_{col}.

Acceptance.

A1 Formal identities/inequalities: first law (5.16), Spohn/Clausius (5.17)–(5.19), Landauer and colored slack, work/free-energy (5.21); **A2** Artifacts in **??**; **A3** Gates G (locality), TQ (CP consistency), S (positivity) verified on the DoVwith reported ζ_{col}.

Part III

Entanglement, Threads, and Device-Independent Constraints

Chapter 6

χ^μ-Threads and Structural Nonseparability

Prelude. We formalize *threads* as conserved, locality–bounded flows that quantify nonseparability in the spectral medium. On a Cauchy slice (Σ, h_{ij}) induced by (M, g_{eff}), a vector density χ^i encodes admissible "entanglement flux" subject to *local norm bounds* fixed by the medium's descriptors (ρ, η, ϕ, T) (H0) and to *locality* (H1). This enables a max–flow/min–cut (MFMC) principle: the thread capacity between regions equals the minimum "cut capacity" of hypersurfaces homologous to those regions. Discretizing Σ yields *thread graphs* with edge capacities and standard network–flow duality. These structures imply quantitative *monogamy* constraints that will feed the device–independent bounds in later chapters.

6.1 Thread Graphs, Cuts, and Monogamy Inequalities

Let Σ be a smooth, compact Cauchy surface with induced metric h_{ij}, area form $\mathrm{d}A = \sqrt{\det h}\ \mathrm{d}^{d-1}x$, and outward unit normal n^i on $\partial\Omega \subset \Sigma$.

6.1.1 Graph models and cuts

Admissible thread fields (continuum model).

Fix a nonnegative *capacity density* $\mathcal{C}(x)$ on Σ (depending on (ρ, η, ϕ, T) via the DoV). A *thread field* is a vector field $\chi^i \in L^2(T\Sigma)$ obeying, distributionally,

$$\nabla_i \chi^i \;=\; 0 \quad \text{on } \Sigma \setminus (\mathcal{W}_{\text{prep}} \cup \mathcal{W}_{\text{meas}}), \qquad \|\chi(x)\|_h \;\le\; \mathcal{C}(x) \quad \text{for a.e. } x \in \Sigma, \tag{6.1}$$

where $\|\chi\|_h = \sqrt{h_{ij}\chi^i\chi^j}$ and $\mathcal{W}_{\text{prep/meas}}$ are localized worldtubes (sources/sinks) compatible with H1 (locality). Denote by $\mathfrak{F}(\mathcal{C})$ the convex set of all such χ.

Flux and capacity between regions.

For disjoint measurable regions $A, B \subset \Sigma$ with smooth boundaries, define the *net flux out of A toward B* carried by χ as

$$\Phi_\chi(A \to B) \;:=\; \int_{\partial A} \chi^i n_i \,\mathrm{d}A \quad \text{with the convention that only threads ending in } B \text{ are counted.} \tag{6.2}$$

The operational *thread capacity* is the MFMC value

$$\mathsf{T}(A:B) \;:=\; \sup_{\chi \in \mathfrak{F}(\mathcal{C})} \Phi_\chi(A \to B). \tag{6.3}$$

Cuts and cut capacity.

Let $\mathcal{S}(A)$ be the family of smooth, oriented hypersurfaces $S \subset \Sigma$ *homologous* to A (i.e. S separates A from B). Define the *cut functional*

$$\mathsf{C}(S) \;:=\; \int_S \mathcal{C}(x)\,\mathrm{d}A(x), \qquad \mathsf{C}_\star(A:B) \;:=\; \inf_{S \in \mathcal{S}(A)} \mathsf{C}(S). \tag{6.4}$$

Theorem 6.1 (Continuum MFMC for threads)**.** *Under H0–H1 and the admissibility constraints* (6.1)*,*

$$\boxed{\mathsf{T}(A:B) \;=\; \mathsf{C}_\star(A:B)}, \qquad \textit{(max–flow = min–cut).} \tag{6.5}$$

Moreover, both the supremum and the infimum are attained whenever $\mathcal{C} \in L^\infty(\Sigma)$ and Σ is compact with piecewise-smooth $\partial A, \partial B$.

Proof sketch. Consider the convex optimization $\max_\chi \int_{\partial A} \chi \cdot n$ subject to (6.1). Introduce a scalar potential $u \in \mathrm{BV}(\Sigma)$ as a Lagrange multiplier for $\nabla \cdot \chi = 0$; the pointwise bound is dualized by the support function of the h-ball of radius $\mathcal{C}$. Fenchel–Rockafellar duality yields the dual problem $\inf_{u:\, u|_A = 1, u|_B = 0} \int_\Sigma \mathcal{C}\, \|\nabla u\|_h \,\mathrm{d}A$, whose minimizers are characteristic functions of cuts S homologous to A. This identifies the dual with (6.4) and proves (6.5). □

Discretization: thread graphs.

Partition Σ into cells $\{V_k\}_{k=1}^N$ by a mesh adapted to the device geometry. For adjacent cells (i, j) let S_{ij} be their common face and set edge capacity $c_{ij} := \int_{S_{ij}} \mathcal{C}\,\mathrm{d}A$. Define the *thread graph* $G = (\mathcal{V}, \mathcal{E}, c)$ with $\mathcal{V} = \{V_k\}$ and $\mathcal{E} = \{(i, j) : S_{ij} \neq \varnothing\}$. A *graph thread flow* is a skew-symmetric function $f : \mathcal{E} \to \mathbb{R}$ such that

$$\sum_{j:(i,j)\in\mathcal{E}} f_{ij} = 0 \quad (\text{Kirchhoff at } i \notin A \cup B), \qquad |f_{ij}| \le c_{ij}.$$

The *graph capacity* is $\mathsf{T}_G(A:B) := \max_f \sum_{i\in A}\sum_j f_{ij}$. Standard network–flow theory yields

$$\mathsf{T}_G(A{:}B) \;=\; \min_{\text{graph cuts } (A|\bar{A})} \sum_{(i,j)\in(A,\bar{A})} c_{ij}, \tag{6.6}$$

and $\mathsf{T}_G \to \mathsf{T}$ as the mesh is refined with consistent capacities c_{ij} (finite–element convergence).

Operational identification.

We take $\mathsf{T}(A{:}B)$ as the *thread-adjusted* bipartite measure for structural nonseparability on Σ. When the POVM coarse–graining (§2.3) respects the partition (A, B), T bounds the achievable CHSH/Tsirelson saturations under the locality and CP gates, to be developed in later chapters.

6.1.2 Monogamy relations

Tripartite monogamy.

Let $A, B, C \subset \Sigma$ be pairwise disjoint. Conservation together with the pointwise bound in (6.1) imply that the total outward flux from A cannot exceed the capacity of any cut separating A from $B \cup C$:

$$\mathsf{T}(A{:}B) + \mathsf{T}(A{:}C) \;\le\; \mathsf{T}(A{:}B\cup C) \;=\; \mathsf{C}_\star(A{:}B\cup C). \tag{6.7}$$

Proof sketch. Superpose admissible maximizers $\chi^{(AB)}$ and $\chi^{(AC)}$ after projecting both into the cone $\|\chi\|_h \le \mathcal{C}$ via convexity; the resulting field carries flux at most $\mathcal{C}$ through every point, so its total flux across any separating cut S is bounded by $\int_S \mathcal{C}\,\mathrm{d}A$. Optimizing over S gives (6.7). □

Polymatroid (submodularity) inequalities.

Define $F(X) := \mathsf{T}(X : \bar{X})$ for measurable $X \subset \Sigma$. Then F is *submodular*:

$$F(X) + F(Y) \;\ge\; F(X\cap Y) + F(X\cup Y), \tag{6.8}$$

and *monotone*: $X \subseteq Y \Rightarrow F(X) \le F(Y)$.

Proof idea. In the dual (cut) picture, $F(X)$ is the minimum of the $\mathcal{C}$-weighted perimeter over cuts homologous to X; weighted perimeter is submodular under union/intersection (BV theory). Duality then transfers submodularity to F. □

Equation (6.8) implies, in particular,

$$\mathsf{T}(A{:}B) + \mathsf{T}(A{:}C) \;\le\; \mathsf{T}(A{:}B\cup C), \qquad \sum_{m=1}^{k} \mathsf{T}(A{:}B_m) \;\le\; \mathsf{T}\big(A{:}\cup_m B_m\big). \tag{6.9}$$

Graph monogamy.

On the thread graph G, let $c(\delta A)$ be the capacity of the edge cut δA. Then the graph analogue reads

$$\mathsf{T}_G(A\!:\!B) + \mathsf{T}_G(A\!:\!C) \;\le\; c(\delta A) \;=\; \min_{\text{graph cuts separating } A} \sum_{e\in\delta A} c_e, \tag{6.10}$$

with equality when the minimal cut is unique and the maximizers are supported on edge–disjoint paths.

Relation to pointer algebra and locality.

If the pointer coarse–graining (§2.3) respects a partition $A \cup B \cup C = \Sigma$ and instruments are localized (H1), then any DI test (CHSH/SDP witnesses) on A and B is constrained by (6.7): attempting to saturate a Bell–type bound between A and B reduces the available capacity for A–C correlations on the same slice. This will be made quantitative in §7.

Acceptance (for this section).

A1 Formal statements/proofs: Theorem 6.1, monogamy (6.7), submodularity (6.8); **A2** Artifacts: mesh builders for thread graphs, max–flow solvers, cut extractors, continuum BV solvers; **A3** Gate compliance: locality (H1) via support in $\mathcal{N}_{L_{\rm loc}}(\mathcal{W})$, positivity (H2/H3) through convex feasibility of $\mathfrak{F}(\mathcal{C})$.

6.2 Thread-adjusted Measures (Negativity/Rel. Entropy): Continuity, Faithfulness

We introduce *thread-adjusted* bipartite measures on a Cauchy slice $\Sigma = A \cup B$ that incorporate the locality/capacity constraints induced by admissible thread fields χ (§6.1.1). The guiding principle is to restrict witnesses/approximants to those that are *implementable* by localized, capacity-bounded operations on $\mathcal{N}_{L_{\rm loc}}(\mathcal{W})$ (Gates $\mathsf{G}, \mathsf{TQ}, \mathsf{S}$). This yields monotones under *thread-LOCC* (local CPTP maps plus classical coordination mediated within the thread budget) and equips the measures with sharp upper bounds in terms of the max–flow/min–cut capacity $\mathsf{T}(A\!:\!B)$.

6.2.1 Adjusted negativity

Standard negativity.

For a bipartite state ρ_{AB} on $\mathcal{H}_A \otimes \mathcal{H}_B$,

$$\mathcal{N}(\rho_{AB}) \;:=\; \tfrac{1}{2}\Big(\|\rho_{AB}^{T_B}\|_1 - 1\Big) \;=\; \sum_{\lambda_i<0} |\lambda_i(\rho_{AB}^{T_B})|,$$

with T_B the partial transpose on B. It admits an SDP representation and a dual in terms of entanglement witnesses $W = W^\dagger$ with $-\mathbb{I} \preceq W^{T_B} \preceq \mathbb{I}$.

Thread-compatible witnesses.

Fix a capacity density $\mathcal{C} \in L^\infty(\Sigma)$ and let $\mathfrak{F}(\mathcal{C})$ be the admissible thread cone (6.1). We say a Hermitian W supported on $A \cup B$ is *thread-compatible* if it can be implemented by a localized meter–coupling whose readout currents are driven by some $\chi \in \mathfrak{F}(\mathcal{C})$ on Σ, in the sense that its operational seminorm obeys the cut bound

$$\|W\|_{\mathrm{cut}(\mathcal{C})} \;:=\; \sup_{S \in \mathcal{S}(A)} \frac{\|\,\mathrm{Tr}_{S^{\mathfrak{C}}}(W)\,\|_\infty}{\displaystyle\int_S \mathcal{C}\,\mathrm{d}A} \;\le\; 1. \tag{6.11}$$

(Here $\mathrm{Tr}_{S^{\mathfrak{C}}}$ contracts directions tangent to the cut S; any equivalent cut–dominated seminorm is acceptable—(6.11) encodes the MFMC restriction that no witness can interrogate more flux than the cut allows.)

Definition (thread-adjusted negativity).

The *thread-adjusted negativity* is

$$\boxed{\mathcal{N}_{\mathcal{C}}(\rho_{AB}) \;:=\; \sup\Big\{ -\tfrac{1}{2}\,\mathrm{Tr}(W\,\rho_{AB}^{T_B}) \;:\; W = W^\dagger,\; -\mathbb{I} \preceq W^{T_B} \preceq \mathbb{I},\;\; \|W\|_{\mathrm{cut}(\mathcal{C})} \le 1 \Big\}.} \tag{6.12}$$

Thus only *thread-implementable* witnesses are admissible in the dual.

Theorem 6.2 (Basic properties of $\mathcal{N}_{\mathcal{C}}$)**.** *Let $\mathsf{T}(A{:}B)$ be the thread capacity (6.3). Then:*

(i) ***Upper/lower bounds.*** $0 \le \mathcal{N}_{\mathcal{C}}(\rho_{AB}) \le \min\{\mathcal{N}(\rho_{AB}),\ \frac{1}{2}\,\mathsf{T}(A{:}B)\}$.

(ii) ***Monotonicity under thread-LOCC.*** *For any CPTP map Φ realizable by local operations on A, B and classical coordination confined to $\mathfrak{F}(\mathcal{C})$, $\mathcal{N}_{\mathcal{C}}(\Phi(\rho)) \le \mathcal{N}_{\mathcal{C}}(\rho)$.*

(iii) ***Faithfulness (on thread-probeable entanglement).*** *If $\mathsf{T}(A{:}B) > 0$ and ρ_{AB} is NPT (i.e. $\rho_{AB}^{T_B} \not\succeq 0$), then $\mathcal{N}_{\mathcal{C}}(\rho_{AB}) > 0$. Conversely, if $\mathsf{T} = 0$, then $\mathcal{N}_{\mathcal{C}} \equiv 0$.*

(iv) ***Continuity.*** *For states ρ, σ on dimension d, if $\|\rho - \sigma\|_1 \le \delta$, then $|\mathcal{N}_{\mathcal{C}}(\rho) - \mathcal{N}_{\mathcal{C}}(\sigma)| \le \frac{1}{2}\,\delta$ and, more sharply, $|\mathcal{N}_{\mathcal{C}}(\rho) - \mathcal{N}_{\mathcal{C}}(\sigma)| \le \frac{1}{2}\,\delta\,\min\{1, \mathsf{T}\}$.*

Proof sketch. (i) The first inequality is by definition. The bound by $\mathcal{N}$ follows since the feasible set in (6.12) is a subset of the standard negativity dual. The capacity bound comes from $|\,\mathrm{Tr}(W\,\rho^{T_B})| \le \|W\|_{\mathrm{cut}}\,\sup_S \int_S \mathcal{C}\,\mathrm{d}A$ by (6.11) and MFMC, yielding $\le \frac{1}{2}\,\mathsf{T}$. (ii) Dual witnesses evolve by the Heisenberg adjoint Φ^*, which preserves the constraints $-\mathbb{I} \preceq (\Phi^* W)^{T_B} \preceq \mathbb{I}$ (CP) and $\|\Phi^* W\|_{\mathrm{cut}} \le \|W\|_{\mathrm{cut}}$ (locality/capacity contraction), so the supremum cannot increase. (iii) If ρ^{T_B} has a negative eigenvalue, a standard optimal witness exists; its localized average can be projected onto the $\|\cdot\|_{\mathrm{cut}} \le 1$ cone without

nullifying the expectation as long as $\mathsf{T} > 0$. (iv) Hölder inequality with $\|W\|_\infty \le 1$ gives the $\frac{1}{2}\,\delta$ bound; the cut seminorm yields the factor $\min\{1, \mathsf{T}\}$. □

Operational remark.

In the tomography-based implementation, the constraint $\|W\|_{\mathrm{cut}} \le 1$ is enforced by sampling W as averages of localized current couplings whose space–time supports are certified to lie within $\mathcal{N}_{L_{\mathrm{loc}}}(\mathcal{W})$ and whose aggregate flux saturates at most one unit of capacity per minimal cut.

6.2.2 Relative-entropy based

Thread-LOCC and thread-separable states.

Let $\mathrm{T-LOCC}_{\mathcal{C}}$ be the set of channels generated by: (i) local CPTP maps on A and B; (ii) classical coordination whose communication and probe patterns are mediated by admissible threads $\mathfrak{F}(\mathcal{C})$ (no side channels across cuts beyond T). Define the *thread-separable* set

$$\mathrm{SEP}_{\mathcal{C}} \;:=\; \overline{\mathrm{conv}}\Big\{\Phi\big(\sum_k p_k\,\rho_k^A \otimes \rho_k^B\big) \;:\; \Phi \in \mathrm{T-LOCC}_{\mathcal{C}},\ \{p_k, \rho_k^A, \rho_k^B\}\ \text{finite}\Big\}. \tag{6.13}$$

When $\mathsf{T} = 0$, $\mathrm{SEP}_{\mathcal{C}}$ reduces to the usual separable set.

Definition (thread-adjusted relative entropy).

$$\boxed{E_R^{\mathrm{thr}}(\rho_{AB}\,;\mathcal{C}) \;:=\; \inf_{\sigma\in\mathrm{SEP}_{\mathcal{C}}} D(\rho_{AB}\,\|\,\sigma), \qquad D(\rho\|\sigma) := \mathrm{Tr}\Big[\rho(\ln\rho - \ln\sigma)\Big].} \tag{6.14}$$

Theorem 6.3 (Properties of E_R^{thr})**.** *For fixed $\mathcal{C}$:*

(i) ***Monotonicity.*** *$E_R^{\mathrm{thr}}(\Phi(\rho)) \le E_R^{\mathrm{thr}}(\rho)$ for all $\Phi \in \mathrm{T-LOCC}_{\mathcal{C}}$.*

(ii) ***Faithfulness.*** *$E_R^{\mathrm{thr}}(\rho) = 0$ iff $\rho \in \overline{\mathrm{SEP}_{\mathcal{C}}}$. If $\mathsf{T} = 0$, this reduces to standard faithfulness w.r.t.* SEP.

(iii) ***Continuity (Alicki–Fannes).*** *If $\|\rho - \sigma\|_1 \le \delta$ on dimension d, then*

$$|E_R^{\mathrm{thr}}(\rho) - E_R^{\mathrm{thr}}(\sigma)| \;\le\; 2\delta\log d + 2h_2(\delta/2),$$

with h_2 the binary entropy.

(iv) ***Pinsker–type lower bound.*** *For the optimizer $\sigma^\star \in \mathrm{SEP}_{\mathcal{C}}$,*

$$E_R^{\mathrm{thr}}(\rho) \;\ge\; \frac{1}{2\ln 2}\,\|\rho - \sigma^\star\|_1^2.$$

(v) ***Capacity upper bound.*** *There exists a constant K (device/calibration dependent) such that $E_R^{\mathrm{thr}}(\rho_{AB}) \;\le\; K\,\mathsf{T}(A\!:\!B)$.*

Proof sketch. (i) Data processing of relative entropy and $\Phi(\mathrm{SEP}_{\mathcal{C}}) \subseteq \mathrm{SEP}_{\mathcal{C}}$. (ii) Direct from the definition and closure. (iii) Standard continuity bound for $D(\cdot\|\cdot)$ with infimum over a fixed convex set. (iv) Pinsker inequality. (v) Construct $\sigma \in \mathrm{SEP}_{\mathcal{C}}$ by coarse–graining across a minimal cut and transporting at most T units of classical/ancilla correlation; bound the divergence by the cut capacity times a local Lipschitz constant K for the implemented channel (on the DoV). □

Relations between $\mathcal{N}_{\mathcal{C}}$ and E_R^{thr}.

Both are monotone under thread-LOCC and admit capacity upper bounds. Using the variational characterization of the trace norm and Pinsker,

$$\mathcal{N}_{\mathcal{C}}(\rho) \;\le\; \tfrac{1}{2}\|\rho^{T_B} - (\sigma^\star)^{T_B}\|_1 \;\le\; \sqrt{\tfrac{\ln 2}{2}\, E_R^{\text{thr}}(\rho)},$$

for the optimizer $\sigma^\star$ in (6.14). Thus a nonzero E_R^{thr} implies a nonzero $\mathcal{N}_{\mathcal{C}}$ and controls its scale, compatibly with capacity bounds.

Acceptance (for this section).

A1 Formal definitions and properties: Theorems 6.2, 6.3; **A2** Artifacts: SDP solvers for $\mathcal{N}_{\mathcal{C}}$ with the cut–seminorm constraint; variational solvers for E_R^{thr} over $\mathrm{SEP}_{\mathcal{C}}$; mesh-based estimators of $\mathsf{T}(A\!:\!B)$; **A3** Gate compliance: locality/capacity auditing for witnesses/channels, CP consistency, and documented DoV.

6.3 Witnesses and Tomography on g_{eff} Backgrounds; Finite-speed Effects

Prelude. All certification and reconstruction tasks in this volume must respect the emergent causal structure of (M, g_{eff}) (H1–H3) and the capacity constraints induced by admissible thread fields (§6.1). Accordingly, (i) entanglement witnesses are chosen from a *local, thread-compatible* cone; (ii) state/channel tomography is posed as a convex program with CP/TP and *cone* (finite-speed) constraints; and (iii) error bars include both statistical noise and a causal slack that quantifies any (residual) leakage outside the g_{eff}–lightcones on the declared DoV.

6.3.1 SDP witnesses

Local, thread-compatible witness cone.

On a Cauchy slice $\Sigma = A \cup B$, a Hermitian operator $W = W^\dagger$ supported on $A \cup B$ is *admissible* if

$$-\mathbb{I} \;\preceq\; W^{T_B} \;\preceq\; \mathbb{I}, \qquad \|W\|_{\mathrm{cut}(\mathcal{C})} \;\le\; 1, \qquad \operatorname{supp} W \subset \mathcal{N}_{L_{\text{loc}}}(\mathcal{W}), \tag{6.15}$$

where the PPT band enforces a standard relaxation, $\|\cdot\|_{\mathrm{cut}(\mathcal{C})}$ is the cut–seminorm (§6.2) encoding the MFMC capacity bound, and the support constraint implements locality on (M, g_{eff}).

Primal SDP (detection).

Given data ρ_{AB} (or a linear model fitted from samples), certify nonseparability by solving

$$\begin{aligned} \text{maximize} \quad & \mathrm{Tr}(W\,\rho_{AB}^{T_B}) \\ \text{subject to} \quad & W \text{ satisfies (6.15).} \end{aligned} \tag{6.16}$$

If the optimum is positive, W detects entanglement. A device-level robust version adds a penalty for colored deviations ζ_{col} (cf. §4.3): maximize $\mathrm{Tr}(W\,\widehat{\rho}^{T_B}) - 2\,\zeta_{\mathrm{col}}\,\|W\|_{\infty}$.

Dual SDP and cut bound.

The Lagrange dual of (6.16) reads

$$\begin{aligned} \text{minimize} \quad & t \\ \text{subject to} \quad & \begin{cases} t\,\mathbb{I} \;\succeq\; X^{T_B} + Y, \\ \|Y\|^{*}_{\mathrm{cut}(\mathcal{C})} \;\leq\; 1, \\ \mathrm{supp}\,X,\ \mathrm{supp}\,Y \subset \mathcal{N}_{L_{\mathrm{loc}}}(\mathcal{W}), \end{cases} \end{aligned} \tag{6.17}$$

where $\|\cdot\|^{*}_{\mathrm{cut}}$ is the dual norm to the cut–seminorm. Weak/strong duality yield the sharp upper bound $\mathrm{Tr}(W\,\rho^{T_B}) \leq \frac{1}{2}\,\mathsf{T}(A{:}B)$ from (6.5) and (6.15).

Finite-speed slack.

If the physical implementation admits a small leakage ϵ_{cone} of the support outside $\mathcal{N}_{L_{\mathrm{loc}}}(\mathcal{W})$ (due to smoothing or discretization), then the certified witness value degrades at most by

$$\Big|\,\mathrm{Tr}(W\,\rho^{T_B}) - \mathrm{Tr}(W_{\mathrm{loc}}\,\rho^{T_B})\Big| \;\leq\; \epsilon_{\mathrm{cone}}\,\|W\|_{\infty}, \tag{6.18}$$

where W_{loc} is the projection of W onto the cone of operators with the prescribed support. Report ϵ_{cone} on the DoVsheet.

Theorem 6.4 (Soundness and capacity upper bound)**.** *Under H1–H3 and* (6.15)*, any optimal value of* (6.16) *satisfies* $0 \leq \mathrm{opt} \leq \frac{1}{2}\,\mathsf{T}(A:B)$*; if* $\mathrm{opt} > 0$ *and* $\epsilon_{\mathrm{cone}} = 0$*, entanglement across* (A, B) *is certified. With nonzero* $\zeta_{\mathrm{col}}, \epsilon_{\mathrm{cone}}$ *the certificate remains valid provided* $\mathrm{opt} > 2\,\zeta_{\mathrm{col}}\|W^{\star}\|_{\infty} + \epsilon_{\mathrm{cone}}\|W^{\star}\|_{\infty}$.

Artifact 6.1: AR: SDP recovery of planted threads

Artifacts: (i) primal/dual SDPs (6.16)–(6.17) with mesh–based support masks; (ii) verifiers for $\|\cdot\|_{\text{cut}}$ via minimal–cut solvers; (iii) simulators that plant flows saturating a chosen cut and recover the corresponding witnesses; (iv) robust variants that include ζ_{col} and ϵ_{cone} penalties.

6.3.2 Tomography under cone constraints

State tomography (cone–CP–TP program).

Let $\{E_m\}$ be localized POVM effects with supports in $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W})$ and outcomes $\widehat{p}_m$ with variances σ_m^2. Estimate ρ by the convex program

$$\begin{aligned} \text{minimize} \quad & \sum_m \frac{(\operatorname{Tr}(E_m\rho) - \widehat{p}_m)^2}{2\sigma_m^2} \;+\; \lambda\, \mathrm{TV}_{\text{cone}}(\rho) \\ \text{subject to} \quad & \rho \succeq 0,\ \operatorname{Tr}\rho = 1,\ \ \operatorname{supp}\rho \subset \mathcal{N}_{L_{\text{loc}}}(\mathcal{W}), \end{aligned} \tag{6.19}$$

where $\mathrm{TV}_{\text{cone}}$ penalizes leakage across the causal mask (e.g. a weighted trace norm of ρ outside the mask).

Channel tomography (Choi picture with cone mask).

For a process Λ, reconstruct its Choi matrix J by

$$\begin{aligned} \text{minimize} \quad & \sum_m \frac{\big(\operatorname{Tr}(F_m J) - \widehat{y}_m\big)^2}{2\sigma_m^2} \;+\; \lambda\, \mathrm{TV}_{\text{cone}}(J) \\ \text{subject to} \quad & J \succeq 0,\ \ \operatorname{Tr}_{\text{out}} J = \mathbb{I}_{\text{in}},\ \ \operatorname{supp} J \subset \mathcal{N}_{L_{\text{loc}}}(\mathcal{W}) \times \mathcal{N}_{L_{\text{loc}}}(\mathcal{W}). \end{aligned} \tag{6.20}$$

Here each F_m represents a prepared–measured probe (GST frame) and the support constraint enforces finite-speed on both input and output legs.

Error guarantees (statistical + causal).

Assume a well-conditioned frame (restricted isometry on the cone mask) with constant κ. Then, with probability $\geq 1-\delta$,

$$\|\widehat{\rho} - \rho^\star\|_1 \;\leq\; C_1 \frac{\sigma}{\sqrt{M}} \;+\; C_2\, \epsilon_{\text{cone}}, \tag{6.21}$$

$$\|\,\widehat{J} - J^\star\|_1 \;\leq\; C_3 \frac{\sigma}{\sqrt{M}} \;+\; C_4\, \epsilon_{\text{cone}}, \tag{6.22}$$

for numerical constants C_i depending on κ and the mask geometry, $\sigma^2 := \max_m \sigma_m^2$, sample count M, and the same leakage ϵ_{cone} as in (6.18). CP/TP are enforced exactly by the SDP constraints; leakage only affects locality.

Device–independent consistency.

Combining (6.22) with the no–signalling theorem (§4.2) one obtains slicewise DI constraints: if two reconstructions $\widehat{J}_A$ and $\widehat{J}_B$ originate from spacelike-separated worldtubes, then the composite adjoint actions commute up to $\mathcal{O}(\epsilon_{\text{cone}})$; any significant noncommutativity indicates a violation of H1/H2 or miscalibration (H5).

Checklist 6.1: PO (T1–T2): Measures / witnesses

- **T1 (Measures).** Compute thread-adjusted negativity and relative-entropy monotones (§6.2) using the cone–SDP estimator $\widehat{\rho}$ or $\widehat{J}$; report capacity bounds.
- **T2 (Witnesses).** Solve (6.16) with cone/mask; certify entanglement when the robust gap exceeds $2\,\zeta_{\text{col}}\|W\|_\infty + \epsilon_{\text{cone}}\|W\|_\infty$.

Chapter 7

Tsirelson Saturation without Superluminal Signalling

Prelude. Bell–CHSH tests on the emergent background (M, g_{eff}) must be run with *localized* instruments supported in spacelike–separated worldtubes $\mathcal{W}_A, \mathcal{W}_B \subset M$ (H1), CP and calibrated channels (H3,H5), and within the declared DoV. Under these conditions the algebras generated by Alice's observables and Bob's observables commute, and the no–signalling theorem of §4.2 applies. In this chapter we (i) formulate the CHSH value as the operator norm of a *Bell operator* that respects the locality gates; (ii) prove the Tsirelson bound $2\sqrt{2}$ by an operator–norm argument that reduces to 2×2 blocks; and (iii) (next subsection) package this as an SDP that returns both upper bounds and saturating state/measurement pairs compatible with the device constraints.

7.1 Operator-norm CHSH; SDP Bounds; Saturating States/Measurements

7.1.1 CHSH via operator norms

Setup and Bell operator.

Let A_0, A_1 act on Alice's Hilbert space $\mathcal{H}_A$ and B_0, B_1 on Bob's $\mathcal{H}_B$. On the DoVwe assume *dichotomic, projective* observables

$$A_x = A_x^\dagger, \quad A_x^2 = \mathbb{I}, \qquad B_y = B_y^\dagger, \quad B_y^2 = \mathbb{I}, \qquad x, y \in \{0, 1\}, \tag{7.1}$$

with *locality*

$$[A_x \otimes \mathbb{I},\ \mathbb{I} \otimes B_y] = 0, \qquad \operatorname{supp}(A_x) \subset \mathcal{N}_{L_{\text{loc}}}(\mathcal{W}_A), \ \ \operatorname{supp}(B_y) \subset \mathcal{N}_{L_{\text{loc}}}(\mathcal{W}_B), \tag{7.2}$$

and $\mathcal{W}_A, \mathcal{W}_B$ spacelike separated in (M, g_{eff}). The *CHSH operator* is

$$\mathcal{B} \;:=\; A_0 \otimes (B_0 + B_1) \;+\; A_1 \otimes (B_0 - B_1). \tag{7.3}$$

For any bipartite state ρ the CHSH value is $\beta(\rho) := \mathrm{Tr}(\rho\,\mathcal{B})$, and its device–independent quantum supremum equals the operator norm: $\sup_\rho \beta(\rho) = \|\mathcal{B}\|$.

Block reduction (Jordan's lemma).

Because A_0, A_1 are reflections ($A_x^2 = \mathbb{I}$), Jordan's lemma yields an orthogonal decomposition of $\mathcal{H}_A$ into 1- and 2-dimensional subspaces invariant under both A_0, A_1; likewise for B_0, B_1 on $\mathcal{H}_B$. Therefore $\mathcal{B}$ is a direct sum of operators acting on 2×2 *qubit* blocks. Its norm is the supremum of the block norms.

Canonical qubit form and exact bound.

On a given 2×2 block we may choose bases so that

$$A_0 = \sigma_z, \qquad A_1 = \cos\alpha\,\sigma_z + \sin\alpha\,\sigma_x, \qquad B_0 = \sigma_z, \qquad B_1 = \cos\beta\,\sigma_z + \sin\beta\,\sigma_x, \tag{7.4}$$

for some $\alpha, \beta \in [0, \pi]$ (local unitaries absorb phases). A direct Pauli calculation gives the identity

$$\mathcal{B}^2 \;=\; 4\,\mathbb{I} \;+\; 4\,\sin\alpha\,\sin\beta\;\sigma_y \otimes \sigma_y, \tag{7.5}$$

whose eigenvalues are $4\Big(1 \pm |\sin\alpha \sin\beta|\Big)$. Hence

$$\|\mathcal{B}\| \;=\; 2\sqrt{1 + |\sin\alpha\,\sin\beta|} \;\le\; 2\sqrt{2}. \tag{7.6}$$

Taking the supremum over blocks produces the global bound.

Theorem 7.1 (Operator–norm Tsirelson bound on (M, g_{eff}))**.** *Under* (7.1)–(7.2) *(hence no–signalling on the* DOV*), the CHSH value satisfies*

$$\sup_{\rho, A_x, B_y} \mathrm{Tr}(\rho\,\mathcal{B}) \;=\; \sup_{A_x, B_y} \|\mathcal{B}\| \;\le\; 2\sqrt{2},$$

with equality iff there exist 2×2 *blocks with* $|\sin\alpha| = |\sin\beta| = 1$ *in* (7.4)*.*

Saturating configuration (example).

On a 2×2 block choose

$$A_0 = \sigma_z, \quad A_1 = \sigma_x, \qquad B_0 = \tfrac{1}{\sqrt{2}}(\sigma_z + \sigma_x), \quad B_1 = \tfrac{1}{\sqrt{2}}(\sigma_z - \sigma_x), \tag{7.7}$$

and the maximally entangled state $|\Phi^+\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$. Then $\langle\Phi^+|\mathcal{B}|\Phi^+\rangle = 2\sqrt{2}$, i.e. the bound is tight. The locality conditions are satisfied since Alice/Bob operators act on spacelike–separated worldtubes $\mathcal{W}_A, \mathcal{W}_B$.

Remarks on robustness (colored noise and finite speed).

If the implemented instruments are subject to colored corrections (§4.3) of size ζ_{col} and a finite-speed leakage ϵ_{cone} in their supports, then the certified value obeys

$$\beta_{\mathrm{cert}} \;\geq\; \mathrm{Tr}(\widehat{\rho}\,\widehat{\mathcal{B}}) \;-\; 2\,\zeta_{\mathrm{col}}\|\widehat{\mathcal{B}}\| \;-\; \epsilon_{\mathrm{cone}}\|\widehat{\mathcal{B}}\|,$$

so saturation requires beating the combined slack. Both terms are controlled by the filter/embedding design of §5.2–5.1.

7.1.2 SDP bounds & saturators

We cast the operator–norm bound into a *finite* semidefinite program of NPA type (level Q_1) that (i) enforces dichotomy and spacelike commutation, (ii) returns the Tsirelson value $2\sqrt{2}$ in the ideal case, and (iii) yields *saturating* states/measurements by Gram factorization when the optimal moment matrix has rank 2.

Words and moment matrix.

Let $\mathcal{W} = \{\mathbb{I}, A_0, A_1, B_0, B_1, A_0B_0, A_0B_1, A_1B_0, A_1B_1\}$ and, for a bipartite state $|\psi\rangle$, define the (Hermitian) moment/Gram matrix

$$\Gamma_{u,v} \;:=\; \langle\psi|\; u^\dagger v \;|\psi\rangle\,, \qquad u, v \in \mathcal{W}. \tag{7.8}$$

By construction $\Gamma \succeq 0$. Algebraic relations from (7.1)–(7.2) translate into *linear* constraints on Γ:

$$\begin{aligned}
&\text{(Id)} \quad \Gamma_{\mathbb{I},\mathbb{I}} = 1;\\
&\text{(dichotomy)} \quad \Gamma_{A_x,A_x} = \Gamma_{B_y,B_y} = 1, \quad \Gamma_{A_x,A_xw} = \Gamma_{\mathbb{I},w},\ \Gamma_{B_y,B_yw} = \Gamma_{\mathbb{I},w}\ (\forall w \in \mathcal{W});\\
&\text{(commutation)} \quad \Gamma_{A_xB_y,\,w} = \Gamma_{B_yA_x,\,w} \ \text{ and } \ \Gamma_{w,\,A_xB_y} = \Gamma_{w,\,B_yA_x}.
\end{aligned} \tag{7.9}$$

The four correlators $\langle A_xB_y\rangle$ appear as entries $\Gamma_{\mathbb{I},A_xB_y} = \Gamma_{A_x,B_y}$.

Primal SDP (NPA-Q_1).

For the CHSH coefficients $c_{xy} \in \{+1, +1, +1, -1\}$ define the linear form

$$\beta(\Gamma) := \sum_{x,y} c_{xy}\,\Gamma_{A_x,B_y}.$$

The device–independent quantum maximum on the DoVis upper-bounded by

$$\begin{aligned}
&\text{maximize} \quad \beta(\Gamma)\\
&\text{subject to} \quad \Gamma \succeq 0, \quad \Gamma \text{ satisfies (7.9)}.
\end{aligned} \tag{7.10}$$

At level Q_1 this SDP is tight for CHSH:

$$\mathrm{opt}\Big((7.10)\Big) = 2\sqrt{2}. \tag{7.11}$$

Dual SDP (Tsirelson certificate).

The dual provides an explicit *sum-of-squares* (SoS) proof of the bound. There exist scalars $\{\lambda\}$ such that

$$2\sqrt{2}\,\mathbb{I} - \mathcal{B} \;=\; \sum_k Z_k^\dagger Z_k$$

with Z_k linear in $\{A_x \otimes \mathbb{I}, \mathbb{I} \otimes B_y\}$; the dual optimum equals $2\sqrt{2}$, furnishing a symbolic certificate of (7.6).

Saturators from Gram factorization.

If an optimal $\Gamma^\star$ has rank $r \le 2$, factorize $\Gamma^\star = V^*V$ with column vectors $\{v_u\}_{u\in\mathcal{W}} \subset \mathbb{R}^r$. On each 2×2 Jordan block, choose a qubit representation such that

$$\langle\psi|A_x \otimes B_y|\psi\rangle \;=\; v_{A_x}\cdot v_{B_y}, \qquad \|v_{A_x}\| = \|v_{B_y}\| = 1.$$

Then define projective observables by $A_x = v_{A_x}\cdot(\sigma_x,\sigma_z)$ and $B_y = v_{B_y}\cdot(\sigma_x,\sigma_z)$ on $\mathbb{C}^2$, and take $|\psi\rangle = |\Phi^+\rangle$; this reproduces the optimal correlators and reaches $2\sqrt{2}$ (recovering (7.7)). When $r > 2$, apply a standard rounding to extract near-saturating qubit blocks (largest singular subspace).

Robust/physical variants (colored and cone constraints).

To certify *implemented* experiments we include slacks for colored noise and finite-speed leakage:

$$\begin{aligned} \text{maximize} \quad & \beta(\Gamma) - \lambda_{\mathrm{col}}\,\zeta_{\mathrm{col}} - \lambda_{\mathrm{cone}}\,\epsilon_{\mathrm{cone}} \\ \text{subject to} \quad & \Gamma \succeq 0,\ (7.9),\ \ |\Gamma_{A_x,B_y} - \widehat{C}_{xy}| \le \delta_{xy}. \end{aligned} \tag{7.12}$$

Here $\widehat{C}_{xy}$ are empirical correlators with confidence radii δ_{xy}; ζ_{col} and ϵ_{cone} are provided by the CP–embedding/filter analysis (§5.2, §5.1). The program returns a *certified* CHSH value

$$\beta_{\mathrm{cert}} \;\le\; \mathrm{opt}\Big((7.12)\Big) \;\le\; 2\sqrt{2} - \lambda_{\mathrm{col}}\,\zeta_{\mathrm{col}} - \lambda_{\mathrm{cone}}\,\epsilon_{\mathrm{cone}}.$$

Finite statistics and DI region.

Concentration bounds of §3.3 yield δ_{xy} at confidence $1-\alpha$ (sub-Gaussian or exact binomial). The feasible set of (7.12) therefore defines a *device–independent* confidence region for the attainable $\{C_{xy}\}$ under the locality/CP gates. Intersecting this region with $\{\beta > 2\}$ certifies nonlocality even with colored slacks.

Summary.

The SDP (7.10) reproduces Tsirelson's bound and, via Gram factorization, produces saturating qubit observables and state compatible with the gates. The robust form (7.12) translates experimental error, colored noise, and finite-speed leakage into certified DI bounds on β.

7.2 Structure Proof: GKSL Locality ⇒ No-signalling beyond Markovian Embeddings

Idea of the proof. We work on a spacetime slab $R = [t_0, t_1] \times \Sigma$ with two spacelike–separated worldtubes $\mathcal{W}_A, \mathcal{W}_B \subset R$ (Gate G). The reduced Heisenberg flow on observables is generated by a *local* GKSL adjoint $\mathcal{L}^*(t)$ (Markov window) or by a *time–local* CP–divisible adjoint $\mathcal{L}^*_t$ obtained from a *local* pseudomode/ancilla dilation (§5.1). We split

$$\mathcal{L}^*(t) \;=\; \mathcal{L}^*_A(t) \;+\; \mathcal{L}^*_B(t) \;+\; \mathcal{L}^*_{\text{rest}}(t), \tag{7.13}$$

where the densities of $\mathcal{L}^*_A$ (resp. $\mathcal{L}^*_B$) are supported in $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W}_A)$ (resp. $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W}_B)$), and $\mathcal{L}^*_{\text{rest}}$ is supported outside both neighborhoods. Microcausality (H2) implies that, as long as $\mathcal{W}_A$ and $\mathcal{W}_B$ remain spacelike, $\mathcal{L}^*_A$ acts trivially on the algebra $\mathcal{A}(\mathcal{W}_B)$ and vice versa.

Causal Heisenberg evolution.

Let $\mathcal{T}$ denote time ordering and write the adjoint propagator as

$$\Lambda^*_{t_1,t_0} \;=\; \mathcal{T}\exp\!\Big(\int_{t_0}^{t_1} \mathcal{L}^*(s)\,\mathrm{d}s\Big). \tag{7.14}$$

Because of (7.13) and locality, for any $O_B \in \mathcal{A}(\mathcal{W}_B)$

$$\frac{\mathrm{d}}{\mathrm{d}t}\,O_B(t) \;=\; \mathcal{L}^*(t)[O_B(t)] \;=\; \Big(\mathcal{L}^*_B(t) + \mathcal{L}^*_{\text{rest}}(t)\Big)[O_B(t)], \qquad O_B(t_0) = O_B, \tag{7.15}$$

i.e. $O_B(t)$ is independent of $\mathcal{L}^*_A$ while A stays spacelike to B.

Lemma 7.2 (Local commutation on spacelike slabs)**.** *Let U_A, U_B be Stinespring unitaries implementing instruments localized in $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W}_A)$ and $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W}_B)$. If the worldtubes are spacelike in R, then for all $O_B \in \mathcal{A}(\mathcal{W}_B)$ and all $t \in [t_0, t_1]$, $[U_A(t), O_B(t)] = 0$ and $\mathcal{L}^*_A(t)[O_B(t)] = 0$.*

Proof sketch. Microcausality (H2) gives [densities in $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W}_A)$, densities in $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W}_B)$] = 0 for spacelike supports. Time ordering preserves this commutation, yielding the two identities. □

Theorem 7.3 (No-signalling from local GKSL structure)**.** *Assume H1–H3 on the declared DoV. Let $\{\mathsf{Instr}^A_i\}$ be any instrument supported in $\mathcal{W}_A$ and let E^B_j be a POVM supported*

in $\mathcal{W}_B$, with $\mathcal{W}_A, \mathcal{W}_B$ spacelike in R. Then for every initial state ρ supported outside $J^-_{g_{\text{eff}}}(\mathcal{W}_A \cup \mathcal{W}_B)$ the marginal statistics in B are independent of the choice of instrument in A:

$$\sum_i \mathrm{Tr}\Big(E_j^B\, \Lambda_{t_1,t_0}\big[\mathsf{Instr}_i^A(\rho)\big]\Big) \;=\; \mathrm{Tr}\Big(E_j^B\, \Lambda_{t_1,t_0}[\rho]\Big) \qquad \forall j. \tag{7.16}$$

Proof sketch. Heisenberg picture with (7.14)–(7.15): by Lemma 7.2, U_A commutes with the Heisenberg evolution of E_j^B, hence the adjoint action of Instr^A cancels in the trace. The remainder $\mathcal{L}^*_{\text{rest}}$ has support outside both neighborhoods, so it cannot mediate influence from $\mathcal{W}_A$ to $\mathcal{W}_B$ within the slab. □

Remarks.

(i) The proof does not require translation invariance nor stationarity; it only uses slicewise locality and microcausality. (ii) If a Lieb–Robinson–type approximation is used instead of strict lightcones, the equalities above become inequalities with exponentially small tails, reported as a cone–leakage ϵ_{cone} (cf. (6.18)).

7.2.1 Beyond-Markov embeddings

We extend Theorem 7.3 to *non-Markovian* evolutions realized by local CP time–local embeddings (§5.1.1).

Embedded model.

Let the system be enlarged with localized pseudomodes/ancillae $\mathcal{A}$ supported in $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W}_A \cup \mathcal{W}_B)$ and let the joint state Ξ_t evolve by a *local* GKSL generator $\mathcal{L}^*_{\text{ext}}(t)$ on system$\otimes\mathcal{A}$:

$$\dot{O}(t) \;=\; \mathcal{L}^*_{\text{ext}}(t)[O(t)], \qquad \Xi_{t_1} \;=\; \Lambda^{\text{ext}}_{t_1,t_0}[\Xi_{t_0}], \qquad \rho_t \;=\; \mathrm{Tr}_{\mathcal{A}}\, \Xi_t.$$

Eliminating $\mathcal{A}$ yields a time–local CP–divisible adjoint $\mathcal{L}^*_t$ on the system (canonical form (5.5)).

Proposition 7.4 (Causal stability under partial trace)**.** *Let O_B be supported in $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W}_B)$. If $\mathcal{L}^*_{\text{ext}}(t)$ is local and microcausal on the slab, then for all $t \in [t_0, t_1]$*

$$\frac{\mathrm{d}}{\mathrm{d}t}\, O_B(t) \;=\; \mathrm{Tr}_{\mathcal{A}}\Big(\mathcal{L}^*_{\text{ext},B}(t)[\widetilde{O}_B(t)]\Big)\,,$$

where $\widetilde{O}_B(t)$ is any lift of $O_B(t)$ to the extended space supported in $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W}_B) \otimes \mathcal{A}$, and $\mathcal{L}^*_{\text{ext},B}$ is the part supported near $\mathcal{W}_B$. In particular, no term supported near $\mathcal{W}_A$ appears on the RHS.

Proof sketch. Same locality argument as (7.15) applied on the extended space, followed by the complete positivity of the partial trace map. □

Theorem 7.5 (No-signalling beyond Markov via local embeddings)**.** *Under H1–H3 and the embedded dynamics above, instruments* $\{\mathsf{Instr}_i^A\}$ *in* $\mathcal{W}_A$ *and POVMs* $\{E_j^B\}$ *in* $\mathcal{W}_B$ *satisfy the marginal independence* (7.16)*. Equivalently, the reduced adjoint evolution on* $\mathcal{A}(\mathcal{W}_B)$ *is independent of the choice of instrument in* A *as long as the worldtubes are spacelike in* R*.*

Proof sketch. Combine Proposition 7.4 with the Stinespring form of instruments on the extended space. Since Instr^A acts trivially on the Heisenberg evolution of E_j^B by locality, the partial trace preserves the independence after elimination of $\mathcal{A}$. □

Approximate embeddings and quantitative slack.

If the colored dynamics is approximated by an R–pseudomode embedding with diamond error ζ_{col} on $[t_0, t_1]$ (cf. (5.7)), then the no–signalling identity acquires an additive slack:

$$\left|\sum_i \mathrm{Tr}\Big(E_j^B\,\Lambda_{t_1,t_0}\big[\mathsf{Instr}_i^A(\rho)\big]\Big) \;-\; \mathrm{Tr}\Big(E_j^B\,\Lambda_{t_1,t_0}[\rho]\Big)\right| \;\le\; 2\,\|E_j^B\|\,\zeta_{\mathrm{col}} \;\le\; 2\,\zeta_{\mathrm{col}}, \tag{7.17}$$

and similarly a cone–leakage term $\epsilon_{\mathrm{cone}}\|E_j^B\|$ if supports are enforced approximately. Both slacks are reported on the DoVsheet and are controllable by the filter design of §5.2.

Consequences for DI tests.

The SDPs of §7.1.2 remain *sound* when augmented with constraints reflecting (7.17) and cone masks; any certified $\beta > 2$ is thus a valid DI witness of nonlocal correlations under the locality and CP gates even beyond the Markov window.

Checklist 7.1: Structural NS proof: what to verify

- Verify the slicewise split (7.13) and supports of all densities.
- Numerically check $[\mathcal{L}_A^*(t),\, \mathcal{A}(\mathcal{W}_B)] = 0$ and vice versa on a basis of local observables.
- For embeddings, audit that the extended generator is local; confirm (7.16) after partial trace up to $(\zeta_{\mathrm{col}}, \epsilon_{\mathrm{cone}})$.

Acceptance.

A1 Formal statements/proofs: Lemma 7.2, Theorems 7.3, 7.5; **A2** Artifacts: locality auditors, slab simulators, embedding builders, NS–slack estimators; **A3** Gate compliance $(\mathsf{G}, \mathsf{TQ}, \mathsf{S})$ within the declared DoV.

7.3 Device-independent Feasibility Regions under Colored Noise; Finite Statistics

Prelude. Given empirical CHSH data collected on spacelike-separated worldtubes $\mathcal{W}_A, \mathcal{W}_B \subset (M, g_{\text{eff}})$ (Gate G), we construct a *device-independent* (DI) feasibility region for correlators and scores that (i) enforces quantum consistency via an NPA (Q_1) moment matrix, (ii) adds *finite-statistics* confidence radii from concentration bounds, and (iii) widens those radii by *colored* and *finite-speed* slacks (ζ_{col}, ϵ_{cone}) derived earlier (§4.3, §5.1). The result is a certified DI region $\mathcal{R}_{1-\alpha}$ at confidence $1-\alpha$.

7.3.1 Concentration bounds

Empirical correlators and win rate.

On n trials with uniformly random inputs $x, y \in \{0,1\}$ and outcomes $a, b \in \{\pm 1\}$, define correlators

$$\widehat{C}_{xy} \;:=\; \frac{1}{n_{xy}} \sum_{t:\, X_t = x, Y_t = y} a_t b_t, \qquad n_{xy} \approx n/4,$$

and the CHSH estimator $\widehat{\beta} = \widehat{C}_{00} + \widehat{C}_{01} + \widehat{C}_{10} - \widehat{C}_{11}$. Equivalently, for the CHSH *game* win indicator $W_t = \mathbf{1}\{a_t b_t = (-1)^{xy}\}$, the empirical win rate $\widehat{p} = \frac{1}{n}\sum_t W_t$ relates to $\widehat{\beta}$ by

$$\widehat{p} \;=\; \tfrac{1}{2} + \tfrac{1}{8}\,\widehat{\beta}, \qquad \widehat{\beta} \;=\; 8\,\widehat{p} - 4. \tag{7.18}$$

Effective sample size under mixing.

Trials may be *weakly dependent* due to colored dynamics. Let the Liouvillian mixing gap on pointer sectors be $\Delta_{\mathcal{L}} > 0$ (H4), and let τ denote the temporal spacing between trials. Choose a blocking length $b \in \mathbb{N}$ comparable to the locking time $b \asymp \lceil c\,(\Delta_{\mathcal{L}}^{-1}/\tau)\rceil$ (constant $c \geq 1$). Define the *effective* sample size $n_{\text{eff}} := \lfloor n/b \rfloor$. All bounds below hold with n replaced by n_{eff} (Azuma–Hoeffding/Freedman martingale blocking). In the i.i.d. Markov window one sets $b = 1$ so that $n_{\text{eff}} = n$.

Coordinate-wise confidence radii for correlators.

For each (x, y) and confidence level α_{xy}:

- **Hoeffding (sub-Gaussian):**

$$\mathbb{P}\Big(\big|\widehat{C}_{xy} - \mathbb{E}[C_{xy}]\big| \;\leq\; \delta^{\text{H}}_{xy}\Big) \;\geq\; 1 - \alpha_{xy}, \qquad \delta^{\text{H}}_{xy} \;:=\; \sqrt{\frac{1}{2n_{\text{eff}}}\log\frac{2}{\alpha_{xy}}}. \tag{7.19}$$

- **Empirical Bernstein (variance-adaptive):** let $\widehat{v}_{xy}$ be the empirical variance of

$\{a_t b_t\}$ on (x, y)-trials. Then

$$\delta_{xy}^{\mathrm{EB}} \;:=\; \sqrt{\frac{2\widehat{v}_{xy}}{n_{\mathrm{eff}}}\log\frac{3}{\alpha_{xy}}} \;+\; \frac{3}{n_{\mathrm{eff}}-1}\log\frac{3}{\alpha_{xy}}, \tag{7.20}$$

guarantees the same tail probability.

Use a union bound $\alpha_{xy} = \alpha/4$ to obtain a simultaneous rectangle for $(\widehat{C}_{00}, \widehat{C}_{01}, \widehat{C}_{10}, \widehat{C}_{11})$.

Direct bound for $\widehat{\beta}$.

Since $\widehat{\beta}$ is an average of n bounded variables in $[-2, 2]$, one can apply Hoeffding/Bennett directly to get

$$\mathbb{P}\Big(\big|\widehat{\beta} - \mathbb{E}[\beta]\big| \;\le\; \Delta_\beta\Big) \;\ge\; 1-\alpha, \qquad \Delta_\beta \;:=\; 2\sqrt{\frac{2}{n_{\mathrm{eff}}}\log\frac{2}{\alpha}}\,, \tag{7.21}$$

or use (7.18) and an exact Clopper–Pearson interval for p.

Colored and cone slacks (structural).

Channel deviations bounded in diamond norm by ζ_{col} and finite-speed leakage ϵ_{cone} produce *additive* slacks in correlators and CHSH scores (cf. (7.17)):

$$\Big|\mathbb{E}[C_{xy}] - C_{xy}^{\mathrm{ref}}\Big| \;\le\; 2\,\zeta_{\mathrm{col}} + \epsilon_{\mathrm{cone}} \quad\Rightarrow\quad \delta^{\mathrm{col+cone}} \;:=\; 2\,\zeta_{\mathrm{col}} + \epsilon_{\mathrm{cone}}. \tag{7.22}$$

These must be *added* to the statistical radii.

DI feasibility region (definition).

Let Γ be an NPA (Q_1) moment matrix (§7.1.2). For chosen per-arm radii δ_{xy} (e.g. $\delta_{xy}^{\mathrm{EB}}$) define

$$\mathcal{R}_{1-\alpha} := \Big\{ C \in [-1,1]^4 \;:\; \exists\, \Gamma \succeq 0 \;\text{ with } Q_1 \text{ constraints},\; |C_{xy} - \widehat{C}_{xy}| \;\le\; \delta_{xy} + \delta^{\mathrm{col+cone}} \;\forall x, y \Big\}. \tag{7.23}$$

Equivalently (score form), intersect the strip $|\beta - \widehat{\beta}| \le \Delta_\beta + \underbrace{4\,\zeta_{\mathrm{col}} + 2\,\epsilon_{\mathrm{cone}}}_{=:\Delta_\beta^{\mathrm{col+cone}}}$ with the Q_1-projected quantum body.

Guarantee.

Under H1–H5 and the mixing/blocking prescription above, the true quantum correlators lie in $\mathcal{R}_{1-\alpha}$ with probability at least $1-\alpha$. Hence any $\beta > 2$ achieved by optimizing over $\mathcal{R}_{1-\alpha}$ is a valid DI certificate of nonlocality at confidence $1-\alpha$ even with colored noise and finite-speed slacks.

Checklist 7.2: PO (C1–C2): DI constraints

- **C1 (Quantum consistency).** Enforce NPA-Q_1 PSD and linear relations for (A_x, B_y); optionally add locality masks for supports.
- **C2 (Statistics + structure).** Use n_{eff} via blocking; choose $\delta_{xy} \in \{\delta^{\text{H}}, \delta^{\text{EB}}\}$ or a CP interval for $\widehat{p}$; add $\delta^{\text{col+cone}}$ from (7.22).

Artifact 7.1: AR: DI regions and concentration scripts

Artifacts: (i) routines computing n_{eff} from $(\Delta_{\mathcal{L}}, \tau)$ and producing $\delta^{\text{H}}, \delta^{\text{EB}}, \Delta_\beta$; (ii) NPA-$Q_1$ SDP builder intersected with the confidence rectangle/strip to output $\mathcal{R}_{1-\alpha}$; (iii) plotting utilities for feasible (C_{xy}) and β intervals; (iv) regression tests that inject colored/cone slacks and recover conservative DI certificates.

Part IV

Master Spectral Principle and Unified Completion

Chapter 8

Master Spectral Principle (MSP): Minimal Axioms and Corollaries

Prelude. The goal of this Part is to compress the logical content of Vols. I–IV into a *minimal* set of axioms—collectively the *Master Spectral Principle* (MSP)—from which the core structural results (Hilbertization, Born-limit, no-signalling under local GKSL/embeddings, pointer algebra, and the geometry–to–couplings map) follow as corollaries on the declared DoV. The axioms are *operational* and *unit-bearing*: every quantity carries physical units, locality is enforced on (M, g_{eff}) via a radius L_{loc}, and all probabilistic statements are tied to calibrated instruments. We then provide independence/redundancy checks that show why each axiom is needed and which statements can be derived without it.

8.1 Minimal Axiom Set Compressing Vols. I–IV; Independence and Redundancy Checks

8.1.1 S0–S4 (minimal)

S0 (Spectral medium & descriptors).

There exists a *spectral medium* with invariant, unit-bearing descriptors

$$(\rho, \eta, \phi, T, \chi^{\mu}) \ \text{ on a smooth manifold } M,$$

and a causal effective metric/cone (M, g_{eff}) on the DoV. The descriptors are bounded/measurable on the relevant worldtubes and fix (i) scales, (ii) admissible *thread capacities* $\mathcal{C}$ on Cauchy slices (cf. Part III), and (iii) the locality radius $L_{\text{loc}} > 0$.

S1 (Locality and Green–hyperbolicity).

All system–meter and matter–gauge couplings are supported in $\mathcal{N}_{L_{\text{loc}}}(\mathcal{W}) \subset (M, g_{\text{eff}})$ (locality gate G). The quadratic Lagrangian density $\mathcal{L}^{(2)}$ of the linearized fields is

Green–hyperbolic; the Peierls bracket on solutions equals the causal CCR/CAR bracket on the Hadamard class, and microcausality holds on spacelike separation.

S2 (Hilbertization and pointer sectors).

Let Sol be the solution space of the Euler–Lagrange equations for $\mathcal{L}^{(2)}$, and ω_{P} the Peierls pre-symplectic form. Quotient by nulls and complete to obtain a one-particle Hilbert space $\mathcal{H}^{(1)}$; second-quantize as appropriate. Pointer sectors are defined by coarse-grained algebras and stabilized by large (ρ, η); there exists a Liouvillian mixing gap $\Delta_{\mathcal{L}} > 0$ on those sectors (H4).

S3 (Localized CP dynamics and instruments).

On the DoVthe reduced dynamics in a worldtube $\mathcal{W}$ is generated either by a local GKSL generator $\mathcal{L}_t$ or by a *local* CP time–local embedding (pseudomodes/reaction coordinates). Instruments are CP and calibrated: there exist Kraus maps $\{\mathsf{Instr}_i\}$ supported in $\mathcal{N}_{L_{\mathrm{loc}}}(\mathcal{W})$ with effects E_i summing to $\mathbb{I}$ (H5).

S4 (Spectral flooding and Born-limit).

Energy-share dynamics between system and meter admits balance kernels compatible with S3. Under the mixing gap and calibration (S2–S3), the *spectral flooding* process reaches a stationary limit ρ_∞ on each pointer sector with outcome weights

$$p_i \;=\; \mathrm{Tr}(\Pi_i\, \rho_\infty),$$

for projectors/effects Π_i of the calibrated instrument. Finite-time deviations obey sub-Gaussian bounds with rates determined by the locking exponent γ and $\Delta_{\mathcal{L}}$ (cf. Chap. 3).

MSP (Master Spectral Principle). S0–S4 hold on the declared DoV. All constructions and bounds in Vol. V are *conditioned* on these axioms and explicitly state their dependence on $(\rho, \eta, \phi, T, \chi^\mu)$, g_{eff}, L_{loc}, and the calibration data.

Remark 8.1 (Immediate corollaries under MSP). (i) **Hilbertization:** S1⇒S2 (pre-symplectic reduction/GNS). (ii) **No-signalling:** S1∧S3 ⇒ local Heisenberg splits and spacelike commutation ⇒ no-signalling (Sections 4.2 and 7.2). (iii) **Born-limit:** S2∧S3∧S4 ⇒ Theorem (Born-limit) with deviation bounds (Sections 3.2 and 3.3). (iv) **Thread bounds:** S0 fixes $\mathcal{C}$; MFMC gives thread capacities that upper-bound DI witnesses (Sections 6.1 and 6.2).

8.1.2 Independence/redundancy table

The table summarizes (a) why each Si is *not* implied by the others (independence via countermodels on the DoV) and (b) which results rely on it.

Axiom	**Independence witness (not implied by others)**	**Used by / forces**
S0	Take flat $(M, \eta_{\mu\nu})$ with well-defined GKSL and Hilbert space (S1–S3 true), but no unit-bearing descriptors or capacity density $\mathcal{C}$. Thread MFMC and DI capacity bounds fail to be meaningful.	Thread capacities $\mathsf{T}(A : B)$; unit-consistent locality scales; mapping to couplings/angles (later).
S1	Choose nonlocal Lindbladian terms (finite Hilbert space, CP instruments allowed). S2–S3–S4 can hold, but spacelike commutation and no-signalling fail.	Microcausality; no-signalling theorems; cone/finite-speed audits.
S2	Consider a purely classical stochastic model with localized CP instruments (S1,S3) and stationary flooding (S4), but no CCR/CAR Hilbertization—quantum structure and pointer superselection cannot be derived.	GNS/Hadamard construction; pointer sectors; spectral positivity checks.
S3	Take unitary but *non-CP* effective maps (e.g. transposition) within a local Hamiltonian field theory (S1,S2). Born-limit formulas (S4) cease to be probabilistically consistent.	GKSL locality; CP time–local embeddings; instruments/POVMs; contractivity tests.
S4	Let a system satisfy S0–S3 but with no mixing gap or flooding balance—then measurement statistics need not converge to $p_i = \mathrm{Tr}(\Pi_i \rho_\infty)$ and deviation bounds fail.	Born-limit theorem and sub-Gaussian deviations; operational identification of p_i.

Redundancy eliminations (mapping Vols. I–IV → S0–S4).

- *Locality postulates* of Vol. I, II ⇒ S1 (one axiom).
- *Hadamard/microcausality & GNS steps* scattered in Vol. I, III are subsumed by S2.
- *Instrument calibration, GKSL derivations, colored embeddings* of Vol. III, IV are compressed into S3.
- *Flooding/Born and finite-time bounds* of Vol. III become S4.
- All *unit/scale* stipulations and thread capacity choices are centralized in S0.

Consistency with gates and DoV.

Each Si carries explicit DoVannotations and gate checks: S1→ G (support/ g_{eff}), S2→ S (positivity/spectral), S3→ TQ (CP/contractivity), S4→ acceptance budgets (gap, calibration).

Proof obligations (PO) spawned by MSP.

- **(M1)** Verify S1⇒ Peierls=CCR/CAR on Hadamard class (microcausality).
- **(M2)** Prove no-signalling under S1∧S3 (Heisenberg locality split).
- **(M3)** Establish Born-limit and deviations under S2–S4 (mixing+calibration).

These POs are discharged in Chaps. 2, 4, 3 and Sections 3.2 and 3.3.

8.2 Corollaries: Emergent g_{eff} and Ehresmann→YM Connections; Anomalies/Inflow

Scope. Under the MSP axioms (S0–S4), two structural consequences follow on the DoV: (i) a *canonical* emergent causal cone/metric g_{eff} together with an Ehresmann connection on internal fibers that reduces to a principal Yang–Mills (YM) connection; (ii) BRST/Ward identities and their possible *anomalous* breakings classified by descent, with inflow cancellation.

8.2.1 Emergent cone/connection

Cone from principal symbol (Hadamard/Green-hyperbolic).

Let $\mathcal{L}^{(2)}$ be the quadratic operator of S1 acting on a field bundle $\pi : \mathcal{E} \to M$. By Green–hyperbolicity its principal symbol has Lorentzian type. Write the (density–weighted) quadratic form

$$\sigma_2\Big(\mathcal{L}^{(2)}\Big)(x,\xi) \;=\; G^{\mu\nu}(x)\,\xi_\mu\xi_\nu \cdot \mathbb{I}_{\mathcal{E}_x} \qquad (\xi \in T_x^*M), \tag{8.1}$$

with signature $(1, d-1)$ on the DoV. Define the pointwise cone

$$\mathcal{C}_x \;:=\; \{\, \xi \neq 0 \;:\; G^{\mu\nu}(x)\,\xi_\mu\xi_\nu \leq 0 \,\}, \qquad \widehat{\mathcal{C}}_x \;:=\; \{\, v \in T_xM \;:\; v\cdot\xi \leq 0 \;\forall \xi \in \mathcal{C}_x \,\}. \tag{8.2}$$

Microcausality and Hadamard propagation (S1) imply that the advanced/retarded fundamental solutions have wavefront sets supported in $\widehat{\mathcal{C}}_x$, so the conformal class $[g_{\text{eff}}]$ is *fixed* by (8.1). Pick a representative g_{eff} (unit-bearing via S0).

Proposition 8.2 (Emergent metric/cone)**.** *On the DoV, $\mathcal{L}^{(2)}$ induces a unique* conformal *Lorentzian metric class $[g_{\text{eff}}]$ whose null cone equals (8.2). The Hadamard parametrix uses*

the squared geodesic distance $\sigma(x, y)$ built from g_{eff}; microcausality coincides with spacelike separation in (M, g_{eff}).

Ehresmann connection on internal fibers.

Let $\pi : \mathcal{F} \to M$ be the internal (matter) fiber bundle whose fibers carry a positive bundle metric k (S0). An *Ehresmann connection* is a horizontal projector $h : T\mathcal{F} \to T\mathcal{F}$ splitting $T\mathcal{F} = \mathsf{H} \oplus \mathsf{V}$ with $\mathsf{V} = \ker d\pi$. Equivalently, a connection 1-form ω valued in the Lie algebra $\mathfrak{iso}(\mathcal{F}, k)$ satisfies $\ker \omega = \mathsf{H}$.

Reduction to a principal YM connection.

Calibration and symmetry on the DoV(S2–S3) select a compact subgroup $G \subset \mathrm{Iso}(\mathcal{F}, k)$ preserving the instrument and pointer data. Then $\mathcal{F}$ reduces to a principal G–bundle $P \to M$ with connection 1-form $A \in \Omega^1(P, \mathfrak{g})$, curvature $F = dA + A \wedge A$, and Bianchi identity $DF = 0$.

Theorem 8.3 (Ehresmann→YM)**.** *Under S0–S3 and locality on (M, g_{eff}), any admissible Ehresmann connection whose horizontal lifts preserve the calibrated pointer structure reduces to a principal G–connection A on $P \to M$. In local trivializations this yields the usual minimal coupling $D_\mu = \partial_\mu + A^a_\mu T_a$ with T_a skew-k–Hermitian, and the Noether currents of §2.3 are the G–covariant fiber charges.*

Spectral stress/current and field equations (on the DoV).

Let $\mathscr{L}_{\text{mat}}(\psi, D\psi; g_{\text{eff}})$ be the effective Lagrangian density and $\mathscr{L}_{\text{YM}} = -\frac{1}{2g^2}\,\mathrm{Tr}(F \wedge \star F)$. Then the Euler–Lagrange equations read

$$\text{(Einstein–like)} \quad \mathsf{G}_{\mu\nu}(g_{\text{eff}}) \;=\; 8\pi\kappa_G \left(T^{\text{mat}}_{\mu\nu} + T^{\text{YM}}_{\mu\nu}\right), \tag{8.3}$$

$$\text{(Yang–Mills)} \qquad D_\mu F^{\mu\nu} \;=\; J^\nu_{\text{mat}}, \qquad J^\nu_{\text{mat}} := \frac{\partial \mathscr{L}_{\text{mat}}}{\partial A_\nu}, \tag{8.4}$$

with $\mathsf{G}_{\mu\nu}$ the effective Einstein tensor for g_{eff} (when a geometric action is included on the DoV). In the purely kinematic case, (8.4) still follows from Noether/Ward identities (below).

Ward/BRST from MSP.

Gauge transformations are vertical flows of $P \to M$; introduce ghost $c = c^a T_a$ and BRST differential s. On fields Φ:

$$sA = -Dc, \qquad sc = -\tfrac{1}{2}[c, c], \qquad s\psi = c \cdot \psi, \qquad s g_{\text{eff}} = 0,$$

and the Slavnov functional $\mathcal{S}(\Gamma) = 0$ encodes the Ward identities for the generating functional Γ when no anomaly is present.

8.2.2 Anomalies and inflow

Consistent anomalies (local BRST cohomology).

Quantum corrections may break BRST at order $\hbar$: $\mathcal{S}(\Gamma) = \hbar\,\mathcal{A} + \mathcal{O}(\hbar^2)$ with $s\mathcal{A} = 0$. Local solutions are classified by the cohomology $H^d_{\mathrm{loc}}(s\,|\,d)$, equivalently by a $(d{+}2)$–form *anomaly polynomial* $\mathcal{I}_{d+2}$ (polynomial in F and the curvature R of g_{eff}) and its descent:

$$\mathcal{I}_{d+2} = \mathrm{d}\mathcal{Q}_{d+1}, \qquad s\,\mathcal{Q}_{d+1} + \mathrm{d}\,\mathcal{Q}^{(1)}_d = 0, \qquad \mathcal{A} = \int_M \mathcal{Q}^{(1)}_d. \tag{8.5}$$

For chiral matter in even $d = 2n$, $\mathcal{I}_{2n+2}$ is the $(2n{+}2)$–form component of $\widehat{A}(R)\,\mathrm{ch}(F)$ (index density).

Wess–Zumino consistency and currents.

The consistent anomaly obeys the WZ condition $\delta_{\alpha_1}\mathcal{A}(\alpha_2) - \delta_{\alpha_2}\mathcal{A}(\alpha_1) = \mathcal{A}([\alpha_1, \alpha_2])$. Equivalently, the covariant divergence of the current satisfies

$$D_\mu J^\mu_{\mathrm{mat}} \;=\; \mathcal{J}_{\mathrm{anom}}(F, R), \tag{8.6}$$

with $\mathcal{J}_{\mathrm{anom}}$ a local polynomial obtained from $\mathcal{I}_{d+2}$.

Inflow mechanism (cancellation by bulk CS).

Let $M = \partial X$ be the boundary of an auxiliary $(d{+}1)$–manifold X. Choose a Chern–Simons form $\mathrm{CS}_{d+1}(A, \Gamma)$ on X with $\mathrm{d}\,\mathrm{CS}_{d+1} = \mathcal{I}_{d+2}$. Then

$$\delta \int_X \mathrm{CS}_{d+1}(A, \Gamma) \;=\; -\int_M \mathcal{Q}^{(1)}_d \;=\; -\,\mathcal{A}, \tag{8.7}$$

so the boundary anomaly is cancelled by bulk inflow. In mixed gauge–gravitational cases Γ is the Levi–Civita connection of g_{eff}.

Operational acceptance on the DoV.

Gate TQ (CP consistency) and the Born limit require that *net* anomalies vanish for the active matter content:

$$\sum_{\mathrm{reps}} \mathcal{I}_{d+2}(F, R) \;=\; 0 \quad \text{in cohomology,} \quad \text{or is cancelled by inflow (8.7).} \tag{8.8}$$

When (8.8) fails, Slavnov identities pick up a nonzero RHS and the GKSL locality/no–signalling proofs remain formal but the physical quotient (§8.3) is obstructed—this is a *kill signature*.

From Ward to YM equations.

In the anomaly–free case, BRST invariance of the effective action implies the Ward identity $D_\mu\langle J^\mu_{\mathrm{mat}}\rangle = 0$. Varying the YM action and using $DF = 0$ recovers (8.4) as the

expectation–value equation of motion. In particular, currents built from the fiber Noether charges (§2.3.1) are G–covariantly conserved on the DoV.

Summary (corollaries under MSP).

- The conformal class $[g_{\text{eff}}]$ and its cone are fixed by the principal symbol of $\mathcal{L}^{(2)}$; the Hadamard machinery and microcausality follow.
- Calibrated horizontal lifts reduce the Ehresmann connection to a principal YM connection A, with standard curvature/Bianchi identities.
- BRST/Ward identities hold; anomalies are classified by $\mathcal{I}_{d+2}$ and can be cancelled by inflow. Acceptance requires (8.8).

8.3 Ward/BRST Structures as MSP Consequences; Physical-Quotient Equivalence

Scope. Under MSP (S0–S4) the gauge structure fixed by the Ehresmann→YM reduction (§8.2.1) implies Slavnov–Taylor (ST) identities and a BRST cohomological description of observables. We develop: (i) the *algebraic* ST/BV machinery on (M, g_{eff}); (ii) renormalization/consistency via the linearized ST operator and local cohomology; and (iii) the *equivalence* between BRST cohomology and the physical (Dirac) quotient, including positivity and the Kugo–Ojima quartet mechanism adapted to pointer sectors.

8.3.1 ST identities

Fields, sources and BRST.

Let $\Phi^I = \{A^a_\mu, \psi, \bar{\psi}, c^a, \bar{c}^a, b^a, \ldots\}$ denote the set of gauge, matter and gauge-fixing fields (including the Nakanishi–Lautrup b and ghost/antighosts $c, \bar{c}$). Introduce anti-fields (Zinn–Justin sources) Φ^*_I coupled to BRST variations $s\Phi^I$ and define the classical gauge-fixed action (on (M, g_{eff}))

$$S_{\text{cl}}[\Phi, \Phi^*] = \int_M \mathrm{d}^d x \sqrt{|g_{\text{eff}}|} \left(\mathscr{L}_{\text{YM}} + \mathscr{L}_{\text{mat}} + s\Psi + \sum_I \Phi^*_I \, s\Phi^I \right), \qquad s^2 = 0. \tag{8.9}$$

Here Ψ is a gauge-fixing fermion; s acts as $sA_\mu = -D_\mu c$, $sc = -\frac{1}{2}[c, c]$, $s\bar{c} = b$, $sb = 0$, $s\psi = c \cdot \psi$, and leaves g_{eff} inert, as in §8.2.1.

Slavnov functional and master equation.

For the vertex functional Γ (generator of 1PI diagrams) define

$$\mathcal{S}(\Gamma) := \int_M \mathrm{d}^d x \sqrt{|g_{\text{eff}}|} \sum_I \frac{\delta \Gamma}{\delta \Phi^*_I} \frac{\delta \Gamma}{\delta \Phi^I} \;\equiv\; \tfrac{1}{2}(\Gamma, \Gamma)_{\text{BV}}. \tag{8.10}$$

ST identity: in the absence of anomalies (cf. §8.2.2),

$$\mathcal{S}(\Gamma) = 0. \tag{8.11}$$

The *linearized* ST operator around Γ is

$$\mathcal{B}_\Gamma X := \int_M \mathrm{d}^d x \sqrt{|g_{\mathrm{eff}}|} \sum_I \left(\frac{\delta\Gamma}{\delta\Phi_I^*} \frac{\delta X}{\delta\Phi^I} + \frac{\delta X}{\delta\Phi_I^*} \frac{\delta\Gamma}{\delta\Phi^I} \right), \qquad \mathcal{B}_\Gamma^2 = 0 \ \text{ if } \ \mathcal{S}(\Gamma) = 0. \tag{8.12}$$

Ward identities and gauge-parameter flows.

Let α denote gauge-fixing parameters entering Ψ. Differentiating (8.11) and using (8.12) yields Nielsen-type identities for gauge-parameter dependence,

$$\frac{\partial\Gamma}{\partial\alpha} = \mathcal{B}_\Gamma\big(\Delta_\alpha\big), \qquad \Rightarrow \qquad \frac{\partial}{\partial\alpha}\Big[\mathcal{O}\Big] = \Big[\mathcal{B}_\Gamma(\cdots)\Big] \ \Rightarrow \ \text{gauge independence of BRST cohomology classes,} \tag{8.13}$$

so S-matrix elements and pointer-sector probabilities computed from physical insertions are α-independent on the DoV.

Local cohomology and counterterms.

Order-by-order renormalization preserves (8.11) iff the *breaking* $\Delta^{(n)}$ at loop order n solves $\mathcal{B}_{\Gamma^{(0)}}\Delta^{(n)} = 0$ and is $\mathcal{B}$-exact, $\Delta^{(n)} = \mathcal{B}_{\Gamma^{(0)}}\Xi^{(n)}$, where $\Gamma^{(0)}$ is the tree functional. Thus counterterms live in $H^0(\mathcal{B}_{\Gamma^{(0)}})$; possible anomalies correspond to $H^1(\mathcal{B}_{\Gamma^{(0)}})$, classified by the descent from the anomaly polynomial (§8.2.2). Under the acceptance condition (8.8) these cohomologies are trivial in the relevant ghost numbers and the ST identity is stable at all orders.

Operators and insertions on (M, g_{eff}).

Local covariant insertions are equivalence classes in the *local* BRST cohomology $H^0_{\mathrm{loc}}(s\,|\,d)$; spacetime derivatives are modded out. Hadamard/Green–hyperbolicity (S1) ensures microlocal control of products, allowing renormalized composite operators compatible with (8.11).

8.3.2 Physical-quotient equivalence

BRST charge and physical space.

Quantize on a Cauchy slice (Σ, h) with Hilbert space $\mathcal{H}$ from S2. The BRST current j^μ_{BRST} is conserved on-shell and defines the (essentially) self-adjoint charge $Q_{\mathrm{BRST}} = \int_\Sigma j^\mu_{\mathrm{BRST}}\, \mathrm{d}\Sigma_\mu$, which is nilpotent, $Q^2_{\mathrm{BRST}} = 0$. The *BRST physical space* is the cohomology

$$\mathcal{H}^{\mathrm{BRST}}_{\mathrm{phys}} \ := \ \ker Q_{\mathrm{BRST}} \big/ \overline{\mathrm{im}\, Q_{\mathrm{BRST}}}. \tag{8.14}$$

Dirac (constraint) quantization.

Let $\mathfrak{G}$ be the group generated by the Gauss constraints on the slice. The *Dirac physical space* is the gauge-invariant subspace $\mathcal{H}_{\text{phys}}^{\text{Dirac}} := \{\, |\Psi\rangle \in \mathcal{H} \,:\, U(g)\,|\Psi\rangle = |\Psi\rangle \;\; \forall g \in \mathfrak{G} \,\}$ modulo null vectors from the indefinite metric (if any).

Theorem 8.4 (Physical-quotient equivalence on the DoV)**.** *Assume S1–S3 and anomaly cancellation* (8.8)*. Then*

$$\mathcal{H}_{\text{phys}}^{\text{BRST}} \;\simeq\; \mathcal{H}_{\text{phys}}^{\text{Dirac}},$$

and expectation values of gauge-invariant, local observables coincide in both descriptions. Moreover, in each pointer sector the Kugo–Ojima quartet mechanism holds: unphysical longitudinal/ghost excitations form BRST doublets whose contribution cancels in physical correlators.

Proof sketch. Construct Q_{BRST} from the Noether charge of the BRST symmetry associated with (8.11). Nilpotency and essential self-adjointness are guaranteed by locality and microlocal spectrum condition (S1–S2) on the DoV. Gauge-invariant states are BRST-closed; BRST-exact states are null in Dirac's inner product. Conversely every Dirac-physical state has a representative in $\ker Q_{\text{BRST}}$. Quotienting by $\overline{\operatorname{im} Q_{\text{BRST}}}$ implements the Gupta–Bleuler/Kugo–Ojima construction. Stability under renormalization follows from the triviality of $H^1(\mathcal{B})$ (no anomaly). □

Pointer sectors and positivity.

In the presence of coarse-grained pointer algebras (S2) the GNS representations are positive; projection to $\mathcal{H}_{\text{phys}}^{\text{BRST}}$ preserves positivity and superselection. Mixing (H4) ensures that long-time limits (Born-limit) computed in either quotient agree with the operational probabilities $p_i = \operatorname{Tr}(\Pi_i \rho_\infty)$.

Equivalence of observable algebras.

Let $\mathfrak{A}_{\text{inv}}$ be the gauge-invariant local algebra on (M, g_{eff}) and $\mathfrak{A}_{\text{phys}}^{\text{BRST}}$ the BRST cohomological algebra of classes $[\mathcal{O}]$ with $s\mathcal{O} = 0$ modulo s-exact terms. Then

$$\mathfrak{A}_{\text{phys}}^{\text{BRST}} \;\simeq\; \mathfrak{A}_{\text{inv}}, \qquad \langle [\mathcal{O}_1] \cdots [\mathcal{O}_n] \rangle = \langle \mathcal{O}_1 \cdots \mathcal{O}_n \rangle_{\text{inv}}, \tag{8.15}$$

compatibly with the Ward identities and with device locality (Gate G).

Operational upshot.

All DI constraints and couplings/mixing predictions that rely on gauge structure can be formulated in either language (Dirac or BRST). The ST identity guarantees their gauge-parameter independence; acceptance requires reporting the cohomology checks and anomaly cancellations used.

Checklist 8.1: PO (M1–M3): Axioms → corollaries

- **M1** (S1⇒S2): verify Peierls=CCR/CAR and microlocal spectrum on (M, g_{eff}) (Hilbertization).
- **M2** (S1∧S3): prove no–signalling/locality of Heisenberg evolution and instruments; ensure ST/BV locality on the slab.
- **M3** (S2–S4): establish Born-limit and show gauge independence of p_i via (8.13); confirm BRST/Dirac equivalence (8.14)–(8.15).

Artifact 8.1: AR: Formal notes and cohomology scripts

Artifacts include: (i) BV/BRST notebooks that build $\mathcal{S}(\Gamma)$ and $\mathcal{B}_\Gamma$, test nilpotency and Nielsen identities; (ii) local cohomology solvers for $H^g_{\text{loc}}(s\,|\,d)$; (iii) constructive Kugo–Ojima routines that identify doublets and project to $\mathcal{H}^{\text{BRST}}_{\text{phys}}$; (iv) regression tests comparing BRST vs Dirac evaluations of observables on toy models on (M, g_{eff}).

Chapter 9

Geometric Determination of Couplings and Mixings

Prelude. Under MSP and the Ehresmann→YM reduction (Section 8.2), the internal fiber $(\mathcal{F}, k)$ comes with a compact symmetry group $G \subset \mathrm{Iso}(\mathcal{F}, k)$ and Lie algebra $\mathfrak{g}$. Geometry on $\mathcal{F}$ fixes an *ad–invariant* inner product on $\mathfrak{g}$; this inner product yields a *gauge–kinetic kernel* whose eigenvalues determine the *bare* gauge couplings. On the declared DoV, thresholding across spectral regimes renormalizes these couplings in a way that preserves positivity and the ST/Ward identities. In this section we build the Lie–metric and extract relations among (g_1, g_2, g_3) (for abelian and nonabelian factors), then state threshold and positivity constraints needed for acceptance gates.

9.1 Lie-metric on $\mathrm{Iso}(\mathcal{F})$ and $g_{1,2,3}$ Relations; Thresholds; Positivity

9.1.1 Metric relations

Fiber–averaged ad–invariant inner product.

Let $X \in \mathfrak{g}$ act on $\mathcal{F}$ via its Killing field K_X (infinitesimal isometry of k). Choose a positive weight density $w(\rho, \eta, \phi, T)$ on $\mathcal{F}$ (unit-bearing, fixed by S0) and define

$$\langle X, Y\rangle_{\mathcal{F}} := \frac{1}{\mathcal{N}_{\mathcal{F}}} \int_{\mathcal{F}} k\Big(K_X, K_Y\Big)\, w \,\mathrm{dvol}_k, \qquad \mathcal{N}_{\mathcal{F}} := \int_{\mathcal{F}} w \,\mathrm{dvol}_k. \tag{9.1}$$

Since $K_{[X,Y]} = [K_X, K_Y]$ and K_X are Killing, (9.1) is *ad–invariant*: $\langle [Z, X], Y\rangle_{\mathcal{F}} + \langle X, [Z, Y]\rangle_{\mathcal{F}} = 0$. Thus $\langle \cdot, \cdot\rangle_{\mathcal{F}}$ is a legitimate Lie–metric on $\mathfrak{g}$ (positive definite for compact G).

Gauge–kinetic kernel and couplings.

Let $\{T_a\}$ be a basis of $\mathfrak{g}$ adapted to the factorization $G = U(1)^m \times \prod_\alpha G_\alpha$ (simple G_α). Define the *gauge–kinetic kernel*

$$\mathsf{G}_{ab} \;:=\; \mathcal{J}\,\langle T_a, T_b\rangle_{\mathcal{F}}, \qquad \mathcal{J} > 0 \text{ (geometric scale, units from S0)}, \tag{9.2}$$

and write the YM term on (M, g_{eff}) as

$$\mathscr{L}_{\text{YM}} \;=\; -\frac{1}{4}\,\mathsf{G}_{ab}\,F^a_{\mu\nu}F^{b\,\mu\nu}. \tag{9.3}$$

By going to the *canonical* kinetic basis (Cholesky or spectral factorization $\mathsf{G} = R^\top R$) the physical couplings are read off from

$$\mathsf{G}_{\text{diag}} \;=\; \text{diag}\left(\frac{1}{g_1^2}, \frac{1}{g_2^2}, \frac{1}{g_3^2}, \dots\right), \qquad \frac{1}{g_a^2} \;=\; \mathcal{J}\,\langle T_a, T_a\rangle_{\mathcal{F}} \quad \text{(no kinetic mixing)}. \tag{9.4}$$

For multiple $U(1)$ factors the abelian block is a positive matrix G_{IJ}; the physical (g_I) and kinetic–mixing angles ε_{IJ} are obtained by diagonalizing G_{IJ}.

Relative coupling relations.

For any two generators T_a, T_b,

$$\frac{g_a^2}{g_b^2} \;=\; \frac{\langle T_b, T_b\rangle_{\mathcal{F}}}{\langle T_a, T_a\rangle_{\mathcal{F}}}, \qquad \Rightarrow \qquad \frac{\alpha_a}{\alpha_b} \;=\; \frac{\langle T_a, T_a\rangle_{\mathcal{F}}}{\langle T_b, T_b\rangle_{\mathcal{F}}}, \quad \alpha_a := \frac{g_a^2}{4\pi}. \tag{9.5}$$

Thus *purely geometric* ratios are predicted at a reference scale μ_{ref} on the DoVbefore RG transport (§10). For a hypercharge generator Y with normalization fixed by $\langle Y, Y\rangle_{\mathcal{F}}$, the weak angle at μ_{ref} obeys

$$\sin^2\theta_W\;(\mu_{\text{ref}}) \;=\; \frac{g_1^2}{g_1^2 + g_2^2} \;=\; \frac{1}{1 + \dfrac{\langle Y, Y\rangle_{\mathcal{F}}}{\langle T^L, T^L\rangle_{\mathcal{F}}}}, \tag{9.6}$$

where T^L denotes any unit generator of the $\mathfrak{su}(2)$ factor in the basis of (9.1). (Care: if Y mixes with other $U(1)$'s, first diagonalize the abelian sub-block.)

Normalization and independence.

If $\langle\cdot,\cdot\rangle_{\mathcal{F}}$ is proportional to the negative Killing form on each simple factor, then $1/g_\alpha^2 \propto \mathcal{J}\,C_\alpha$ with C_α the quadratic Casimir normalization on $\mathfrak{g}_\alpha$. The choice of w and $\mathcal{J}$ sets the overall scale but cancels in (9.5)–(9.6); hence ratios are *independent* of that normalization (MSP prediction at μ_{ref}).

Example (abelian mixing). For two $U(1)$ generators Y_1, Y_2 the abelian kernel is $\mathsf{G}_{IJ} = \mathcal{J}\langle Y_I, Y_J\rangle_{\mathcal{F}}$ (a 2×2 PSD matrix). Physical couplings are given by the eigenvalues of

G and the mixing angle by the orthogonal diagonalizer; (9.6) uses the Y direction that couples to matter after this rotation.

9.1.2 Threshold/positivity

One–loop transport with geometric boundary conditions.

Choose a reference scale μ_{ref} within the DoVwhere the geometric kernel (9.2) is calibrated. For $\mu < \mu_{\text{ref}}$, the (scheme–dependent) one–loop running reads

$$\alpha_a^{-1}(\mu) \;=\; \underbrace{\alpha_a^{-1}(\mu_{\text{ref}})}_{\text{geometry}} \;-\; \frac{b_a}{2\pi}\,\ln\frac{\mu}{\mu_{\text{ref}}} \;-\; \sum_{t\in\mathcal{T}} \frac{\Delta b_a^{(t)}}{2\pi}\,\ln\frac{\mu}{M_t} \;-\; \delta_a^{\text{fin}}, \tag{9.7}$$

with b_a the light–spectrum beta coefficients, thresholds M_t for heavy modes/pseudomodes indexed by t (DoV boundaries), and finite matching δ_a^{fin} extracted from the CP embedding of Chap. 5. Equation (9.7) preserves the *geometric* relations at μ_{ref} and propagates them to experimental scales.

Positivity and kinetic–mixing constraints.

The gauge–kinetic kernel must remain positive definite on the DoV:

$$\mathsf{G}(\mu) \;\succ\; 0 \qquad\Longleftrightarrow\qquad v^\top \mathsf{G}(\mu)\, v \;>\; 0 \quad \forall\, 0 \neq v \in \mathbb{R}^{\dim \mathfrak{g}}, \tag{9.8}$$

and, for abelian subsectors,

$$\mathsf{G}_{U(1)^m}(\mu) \;\succeq\; 0, \qquad \|\mathsf{G}_{U(1)^m}^{-1/2}\,\varepsilon\,\mathsf{G}_{U(1)^m}^{-1/2}\| \;<\; 1 \;\;\text{(bounded kinetic mixing)}. \tag{9.9}$$

These constraints ensure no ghostlike modes and well–posed Cauchy problems for the gauge sector on (M, g_{eff}). They are enforced numerically as LMIs in the predictive pipeline (cf. Sections 5.2 and 5.3).

Threshold consistency (DoV boundaries).

Crossing a DoV boundary (appearance/disappearance of modes) changes $\{b_a\}$ by $\Delta b_a^{(t)}$ in (9.7). MSP requires:

$$\lim_{\epsilon\downarrow 0}\; \alpha_a^{-1}(M_t - \epsilon) \;=\; \alpha_a^{-1}(M_t + \epsilon) \;+\; \delta_a^{\text{fin}}(M_t), \tag{9.10}$$

with $\delta_a^{\text{fin}}(M_t)$ computed from the CP embedding and respecting the ST identities (no gauge–symmetry violation at matching). In multi–$U(1)$ sectors, matching is done at the kernel level (G_{IJ}) before diagonalization.

Geometric bounds and sum rules.

If the matter content fills complete representations of a larger isometry group $\tilde{G} \supset G$ near μ_{ref}, then the fiber inner product descends from $\tilde{\mathfrak{g}}$ and imposes *sum rules* among $1/g_a^2$ via Casimirs of $\tilde{\mathfrak{g}}$. Such relations survive RG flow modulo calculable threshold/finite corrections and provide falsifiable envelopes in Chap. 10.

Acceptance.

Report $(\mathsf{G}_{ab}(\mu_{\text{ref}}), \mathcal{J})$, the geometric ratios (9.5)–(9.6), and the RG/threshold map (9.7); verify positivity (9.8)–(9.9) on the DoVand ST–consistent matching (9.10). These constitute the inputs to the prediction module for $(g_{1,2,3})$, $\sin^2\theta_W$, and mixing textures in Chapter 10.

9.2 Weinberg Angle from Fiber Geometry: Prediction with Uncertainties

Scope. We derive a *geometric* prediction for the weak mixing angle $\sin^2\theta_W$ at a reference scale μ_{ref} directly from the fiber Lie–metric (§9.1.1), including abelian kinetic mixing, and propagate it to the electroweak scale M_Z with quantified uncertainties from geometry, thresholds, and running. Throughout, all statements are on the declared DoVand respect the gates (locality/CP/positivity).

9.2.1 Reference-scale prediction

Set-up and conventions.

Let the internal symmetry be $G = U(1)^m \times \prod_\alpha G_\alpha$ with $\mathfrak{g} = \text{Lie}(G)$. The gauge–kinetic kernel from the Lie–metric (§9.1.1) is

$$\mathsf{G}_{AB} \;=\; \mathcal{J}\,\langle T_A, T_B\rangle_{\mathcal{F}} \qquad (A, B \in \{U(1)^m,\ \text{simple factors}\}), \tag{9.11}$$

positive definite on the DoV. For a simple factor (e.g. $\text{SU}(2)_L$) one has $1/g_2^2 = \mathcal{J}\,\langle T^L, T^L\rangle_{\mathcal{F}}$. For an abelian sector with possible kinetic mixing, write the abelian block as the $m \times m$ SPD matrix G_{IJ} in a basis $\{Y_I\}$ of $U(1)^m$.

Hypercharge coupling with abelian mixing.

Let $q_Y \in \mathbb{R}^m$ be the unit vector selecting the hypercharge direction in the matter charge space (i.e. the covariant derivative couples as $D_\mu = \partial_\mu + \sum_I g_I\, Y_I A_\mu^I$ and physical hypercharge is the combination $Y = q_Y^I Y_I$). After canonical normalization of abelian

kinetics, the physical hypercharge coupling at $\mu_{\rm ref}$ is

$$\boxed{\frac{1}{g_Y^2(\mu_{\rm ref})} \;=\; q_Y^\top\, \mathsf{G}_{U(1)}^{-1}(\mu_{\rm ref})\, q_Y, \qquad \mathsf{G}_{U(1)} := \Big[\mathsf{G}_{IJ}\Big]_{I,J=1}^{m},} \tag{9.12}$$

which reduces to $1/g_Y^2 = \mathsf{G}_{YY}$ when $m = 1$ and no mixing.

Geometric formula for $\sin^2\theta_W$ at $\mu_{\rm ref}$.

With $1/g_2^2 = \mathcal{J}\,\langle T^L, T^L\rangle_{\mathcal{F}} =: v$ and $1/g_Y^2 =: u = q_Y^\top \mathsf{G}_{U(1)}^{-1} q_Y$, one finds the *exact* reference-scale relation

$$\boxed{\sin^2\theta_W(\mu_{\rm ref}) \;=\; \frac{g_Y^2}{g_Y^2+g_2^2} \;=\; \frac{v}{u+v}\,, \qquad u = \; q_Y^\top \mathsf{G}_{U(1)}^{-1} q_Y, \quad v = \; \mathcal{J}\,\langle T^L, T^L\rangle_{\mathcal{F}}.} \tag{9.13}$$

Normalization note. If one compares to conventions using a rescaled hypercharge (e.g. GUT normalization), include the known factor inside q_Y; the formula (9.13) is invariant under simultaneous rescalings of Y and its coupling and thus convention-safe.

Uncertainty at $\mu_{\rm ref}$ (geometric + abelian-mixing).

Let the geometric inputs be $\Theta = \{\mathsf{G}_{U(1)}, v\}$ with covariance $\mathrm{Cov}(\Theta)$. From $s^2 := \sin^2\theta_W = v/(u+v)$ one has

$$\frac{\partial s^2}{\partial u} \;=\; -\frac{v}{(u+v)^2}, \qquad \frac{\partial s^2}{\partial v} \;=\; \frac{u}{(u+v)^2}. \tag{9.14}$$

Since $u = q_Y^\top \mathsf{G}_{U(1)}^{-1} q_Y$, a small variation $\delta\mathsf{G}$ induces

$$\delta u \;=\; -\,q_Y^\top\, \mathsf{G}_{U(1)}^{-1}\,\delta\mathsf{G}\,\mathsf{G}_{U(1)}^{-1}\, q_Y \;=\; -\,\mathrm{Tr}\Big(\mathsf{W}\,\delta\mathsf{G}\Big), \qquad \mathsf{W} := \mathsf{G}_{U(1)}^{-1} q_Y q_Y^\top \mathsf{G}_{U(1)}^{-1}. \tag{9.15}$$

Therefore, to leading order,

$$\mathrm{Var}\Big[s^2(\mu_{\rm ref})\Big] \;\approx\; \Big(\partial_u s^2\Big)^2\,\mathrm{Var}[u] + \Big(\partial_v s^2\Big)^2\,\mathrm{Var}[v] + 2\,(\partial_u s^2)(\partial_v s^2)\,\mathrm{Cov}[u,v], \tag{9.16}$$

with $\mathrm{Var}[u] = \mathrm{Tr}(\mathsf{W}\,\mathrm{Cov}[\mathsf{G}]\,\mathsf{W})$. This captures both geometric uncertainties and abelian-mixing calibration errors.

Acceptance at $\mu_{\rm ref}$.

Report $\Big(\mathsf{G}_{U(1)}(\mu_{\rm ref}), v(\mu_{\rm ref})\Big)$, the central value $s^2(\mu_{\rm ref})$ from (9.13), and its uncertainty from (9.16); verify positivity $\mathsf{G}_{U(1)} \succ 0$ and simple-factor positivity $v > 0$ (Gate S).

9.2.2 RG transport to M_Z

Running with thresholds (matrix form for $U(1)^m$).

Between $\mu_{\rm ref}$ and M_Z the gauge sector runs with beta coefficients $\{b_2,\, \mathbf{B}_{U(1)}\}$ (the latter an $m \times m$ matrix for abelian mixing). In terms of inverse couplings:

$$\alpha_2^{-1}(\mu) \;=\; \alpha_2^{-1}(\mu_{\rm ref}) \;-\; \frac{b_2}{2\pi}\ln\frac{\mu}{\mu_{\rm ref}} \;-\; \sum_t \frac{\Delta b_2^{(t)}}{2\pi}\ln\frac{\mu}{M_t} \;-\; \delta_2^{\rm fin}, \tag{9.17}$$

$$\mathsf{G}_{U(1)}^{-1}(\mu) \;=\; \mathsf{G}_{U(1)}^{-1}(\mu_{\rm ref}) \;-\; \frac{\mathbf{B}_{U(1)}}{2\pi}\ln\frac{\mu}{\mu_{\rm ref}} \;-\; \sum_t \frac{\Delta\mathbf{B}_{U(1)}^{(t)}}{2\pi}\ln\frac{\mu}{M_t} \;-\; \Delta^{\rm fin}, \tag{9.18}$$

with $\Delta b^{(t)}$ (resp. $\Delta\mathbf{B}^{(t)}$) encoding heavy-mode thresholds at M_t (DoV boundaries) and $\delta^{\rm fin}, \Delta^{\rm fin}$ finite matching from the CP embedding (§5). Evaluate the *physical* hypercharge inverse coupling at μ by

$$u(\mu) \;:=\; \frac{1}{g_Y^2(\mu)} \;=\; q_Y^\top\, \mathsf{G}_{U(1)}^{-1}(\mu)\, q_Y, \qquad v(\mu) \;:=\; \frac{1}{g_2^2(\mu)}. \tag{9.19}$$

Predicted $\sin^2\theta_W$ at M_Z.

Insert (9.19) into the exact identity

$$\boxed{\sin^2\theta_W(M_Z) \;=\; \frac{v(M_Z)}{u(M_Z)+v(M_Z)}\,,} \tag{9.20}$$

with $u(M_Z)$ obtained from the *matrix* evolution (9.18) even when $m > 1$.

Uncertainty propagation (running + thresholds).

Let $\Xi = \Big(\alpha_2^{-1}(\mu_{\rm ref}), \mathsf{G}_{U(1)}^{-1}(\mu_{\rm ref}), b_2, \mathbf{B}_{U(1)}, \{M_t\}, \delta_2^{\rm fin}, \Delta^{\rm fin}\Big)$. Linearizing (9.17)–(9.18) and using (9.14) with $u \to u(M_Z)$, $v \to v(M_Z)$ yields

$$\mathrm{Var}\Big[\sin^2\theta_W(M_Z)\Big] \;\approx\; (\nabla s^2)^\top\, \mathrm{Cov}\Big(u(M_Z), v(M_Z)\Big)\,(\nabla s^2), \qquad \nabla s^2 = \begin{pmatrix}\partial_u s^2\\ \partial_v s^2\end{pmatrix}, \tag{9.21}$$

where $\mathrm{Cov}(u(M_Z), v(M_Z))$ follows from the affine dependence on Ξ in (9.17)–(9.18). Threshold uncertainties enter through $\mathrm{Var}[\ln(\mu/M_t)]$ and the matching covariances $\mathrm{Cov}(\delta_2^{\rm fin}, \Delta^{\rm fin})$.

Checks and acceptance.

(i) *Positivity along the flow:* enforce $\mathsf{G}_{U(1)}(\mu) \succ 0$, $v(\mu) > 0$ on $\mu \in [M_Z, \mu_{\rm ref}]$ (LMIs; Gate S). (ii) *ST/Ward consistency:* thresholds/matching must preserve gauge identities (Gate TQ). (iii) *Locality:* running inputs and threshold placements are valid only within the declared DoV(Gate G).

Artifact 9.1: AR: RG notebooks and uncertainty budgets

Artifacts include: (i) routines to construct $\mathsf{G}_{U(1)}$ and v at $\mu_{\rm ref}$ from fiber geometry; (ii) matrix RG evolution for abelian mixing, with threshold/matching support; (iii) automatic evaluation of $\sin^2\theta_W(\mu_{\rm ref})$ and $\sin^2\theta_W(M_Z)$ using (9.13) and (9.20); (iv) uncertainty propagation per (9.16), (9.21), reporting geometric vs running vs threshold components.

9.3 Overlap-driven Yukawa Textures; CKM/PMNS Sum Rules

Prelude. Within MSP, the internal fiber $(\mathcal{F}, k)$ equipped with a compact isometry group $G \subset \mathrm{Iso}(\mathcal{F}, k)$ not only fixes the gauge–kinetic kernel (§9.1.1) but also constrains the *Yukawa sector*. Matter fields carry profiles on $\mathcal{F}$ and couple to a (possibly distributed) Higgs background; fermion masses arise from *overlap integrals* on $\mathcal{F}$ with unit-bearing weight w (S0). This section develops the geometric construction of Yukawa *textures* at a reference scale $\mu_{\rm ref}$ on the declared DoV, including canonical normalization of kinetic terms, selection rules from residual fiber symmetry, and low-rank/perturbative structures that naturally generate hierarchies.

9.3.1 Overlap textures

Fields, profiles, and kinetic norms on $\mathcal{F}$.

Let $\psi_L^{(i)}(x,y)$ and $\psi_R^{(j)}(x,y)$ denote left-/right-handed Weyl components (flavor indices i,j) with $x \in M$ and $y \in \mathcal{F}$. Factorize along $\mathcal{F}$ by smooth profiles $f_L^{(i)}(y)$, $f_R^{(j)}(y) \in \Gamma(\mathcal{F})$ and x-dependent spinors:

$$\psi_L^{(i)}(x,y) = f_L^{(i)}(y)\,\Psi_L^{(i)}(x), \qquad \psi_R^{(j)}(x,y) = f_R^{(j)}(y)\,\Psi_R^{(j)}(x).$$

The fiber-induced kinetic Gram matrices at $\mu_{\rm ref}$ are

$$\big(G_L\big)_{ij} := \int_{\mathcal{F}} \big\langle f_L^{(i)}, f_L^{(j)} \big\rangle_k \, w \,\mathrm{dvol}_k, \qquad \big(G_R\big)_{ij} := \int_{\mathcal{F}} \big\langle f_R^{(i)}, f_R^{(j)} \big\rangle_k \, w \,\mathrm{dvol}_k, \tag{9.22}$$

which are positive definite on the DoV(Gate S).

Bare overlap Yukawas and canonical normalization.

Let $\mathcal{Y}(y)$ be a fiber endomorphism encoding the local Yukawa density and $h(y)$ the fiber profile of the Higgs condensate (or a weighted sum of modes). The *bare* Yukawa matrix at $\mu_{\rm ref}$ is

$$Y_{ij}^{(\rm bare)} := \int_{\mathcal{F}} \big\langle f_L^{(i)}(y),\ \mathcal{Y}(y)\, f_R^{(j)}(y) \big\rangle_k \, h(y)\, w(y)\,\mathrm{dvol}_k(y). \tag{9.23}$$

To read physical couplings, canonically normalize fermions by $\Psi_L \to G_L^{-1/2}\Psi_L$, $\Psi_R \to G_R^{-1/2}\Psi_R$, which yields the *canonically normalized* Yukawa matrix

$$\boxed{\widehat{Y} \;:=\; G_L^{-1/2}\, Y^{(\mathrm{bare})}\, G_R^{-1/2}.} \tag{9.24}$$

All predictions below use $\widehat{Y}(\mu_{\mathrm{ref}})$; RG transport to experimental scales is deferred to §10.

Mode expansion and low-rank structure.

Expand the fiber Yukawa density against a finite set of orthonormal modes $\{\Xi_m(y)\}_{m=1}^r$ (e.g. eigenmodes of a G-covariant Laplacian):

$$\mathcal{Y}(y)\, h(y) = \sum_{m=1}^{r} y_m\, \Xi_m(y), \qquad \int_{\mathcal{F}} \langle \Xi_m, \Xi_{m'} \rangle_k\, w\, \mathrm{dvol}_k = \delta_{mm'}.$$

Then (9.23) factorizes as a *finite sum of rank-one* pieces:

$$Y^{(\mathrm{bare})} \;=\; \sum_{m=1}^{r} y_m\, u^{(m)} v^{(m)\top}, \quad u_i^{(m)} := \int_{\mathcal{F}} \langle f_L^{(i)}, \Xi_m \rangle_k\, w\, \mathrm{dvol}_k, \;\; v_j^{(m)} := \int_{\mathcal{F}} \langle \Xi_m, f_R^{(j)} \rangle_k\, w\, \mathrm{dvol}_k. \tag{9.25}$$

If one mode dominates, Y is approximately rank one, generating a heavy third generation with hierarchical subleading masses from smaller $|y_m|$ or suppressed overlaps—this provides a *geometric* origin for texture hierarchies.

Selection rules from residual fiber symmetry.

Let $H \subseteq G$ be the residual symmetry on the DoV. If $f_L^{(i)} \in \mathrm{Rep}_{\lambda_i}$, $f_R^{(j)} \in \mathrm{Rep}_{\rho_j}$ and $\Xi_m \in \mathrm{Rep}_{\sigma_m}$, then Schur orthogonality implies

$$Y_{ij}^{(\mathrm{bare})} = 0 \quad \text{unless} \quad \mathrm{Rep}_{\lambda_i} \otimes \mathrm{Rep}_{\sigma_m} \otimes \mathrm{Rep}_{\rho_j} \text{ contains the trivial rep.} \tag{9.26}$$

Thus *texture zeros* are symmetry-protected and robust under small deformations that preserve H.

Localized profiles and exponential hierarchies (Gaussian model).

Suppose (for intuition) that left/right profiles are localized Gaussians on $\mathcal{F} = \mathbb{R}^d$ (or a patch) with widths $\Sigma_{L,R}$ and centers y_{L_i}, y_{R_j}, and that $\mathcal{Y}h$ is slowly varying. Then

$$Y_{ij}^{(\mathrm{bare})} \;\propto\; \exp\Big[-\tfrac{1}{2}\big(y_{L_i} - y_{R_j}\big)^{\top} \Sigma^{-1} \big(y_{L_i} - y_{R_j}\big)\Big], \qquad \Sigma := \Sigma_L + \Sigma_R, \tag{9.27}$$

so small separations yield large entries (heavy families), while distant pairs are exponentially suppressed (light families). Canonical normalization rescales rows/columns by $G_{L,R}^{-1/2}$ but preserves the *pattern* of hierarchies.

Biunitary diagonalization and texture parameters.

At μ_{ref} write the singular value decomposition

$$\widehat{Y} \;=\; U_L \,\mathrm{diag}(y_1, y_2, y_3)\, U_R^\dagger, \qquad y_1 \le y_2 \le y_3, \tag{9.28}$$

with $U_{L,R}$ unitary. The mass eigenvalues are $m_i(\mu_{\text{ref}}) = y_i\, v_H/\sqrt{2}$ (for Higgs v.e.v. v_H on the DoV). For two fermion sectors (e.g. up/down quarks), the CKM matrix arises as $V_{\text{CKM}} = U_L^{(u)\dagger} U_L^{(d)}$; similarly PMNS from lepton sectors. (Explicit CKM/PMNS sum rules are developed in §9.3.2.)

Small-parameter expansion (texture bookkeeping).

Let $\varepsilon \ll 1$ bound off-diagonal overlaps in a basis where one mode dominates, so that $\widehat{Y} = y_3 \left(e_3 e_3^\top + \mathcal{O}(\varepsilon)\right)$. Then

$$\frac{y_2}{y_3} = \mathcal{O}(\varepsilon), \quad \frac{y_1}{y_2} = \mathcal{O}(\varepsilon), \quad \theta_{12}, \theta_{23}, \theta_{13} = \mathcal{O}(\varepsilon) \text{ (CKM-like)},$$

or, if Ξ_m carry phases or near-degeneracies, large mixings (PMNS-like) can occur even with hierarchical y_i.

Positivity, units, and acceptance.

The Gram matrices $G_{L,R} \succ 0$ (Gate S) ensure well-posed canonical normalization. All fiber integrals carry units set by w (S0); Y is rendered dimensionless after factoring the Higgs normalization into h. The DoVmust specify the set $\{\Xi_m\}$ and the calibration of w so that (9.25) is reproducible (Gate TQ).

Deliverables (for Y1–Y2).

(i) Construction of $G_{L,R}$ and $\widehat{Y}$ from given profiles and modes; (ii) symmetry audit implementing (9.26); (iii) extraction of texture parameters $(y_i, U_{L,R})$ at μ_{ref}; (iv) export to the RG notebooks for transport and uncertainty budgets.

9.3.2 CKM/PMNS relations

Set-up and notation.

At the reference scale μ_{ref} let the canonically normalized Yukawas from (9.24) be $\widehat{Y}_u, \widehat{Y}_d$ (quarks) and $\widehat{Y}_e, \widehat{Y}_\nu$ (leptons). Define the Hermitian left–invariant matrices

$$H_u := \widehat{Y}_u \widehat{Y}_u^\dagger, \qquad H_d := \widehat{Y}_d \widehat{Y}_d^\dagger, \qquad H_e := \widehat{Y}_e \widehat{Y}_e^\dagger, \tag{9.29}$$

with spectral decompositions $H_u = U_L^{(u)} \mathrm{diag}(y_u^2, y_c^2, y_t^2)\, U_L^{(u)\dagger}$, etc. Then

$$V_{\text{CKM}} \;:=\; U_L^{(u)\dagger} U_L^{(d)}, \qquad U_{\text{PMNS}} \;:=\; U_L^{(e)\dagger} U_\nu, \tag{9.30}$$

where U_ν diagonalizes the (Majorana) neutrino mass matrix M_ν (see below). Angles/phases follow the PDG parameterization.

Geometric content via overlap modes.

Using the overlap expansion (9.25), the *left* structures read

$$H_u = \sum_m \lambda_m^{(u)}\, a_m^{(u)} a_m^{(u)\dagger}, \qquad H_d = \sum_n \lambda_n^{(d)}\, a_n^{(d)} a_n^{(d)\dagger}, \tag{9.31}$$

with $\lambda_m^{(u)} = |y_m^{(u)}|^2$, and left overlap vectors $a_m^{(u)} := G_L^{-1/2} u^{(m)}$ (quark doublets), likewise for d. Rank–one dominance corresponds to one large λ_3 with $a_3^{(u)}, a_3^{(d)}$ approximately aligned.

Small–angle formulae (quarks).

In the basis where H_d is diagonal ("down basis"), denote $\Delta_{ij}^{(u)} := y_{u,j}^2 - y_{u,i}^2$ and the off–diagonals of H_u by $(H_u)_{ij}$. First–order perturbation theory gives

$$\theta_{ij}^{\mathrm{CKM}} \simeq \left|\frac{(H_u)_{ij}}{\Delta_{ij}^{(u)}}\right| \quad \text{(down basis)}, \qquad \theta_{ij}^{\mathrm{CKM}} \simeq \left|\frac{(H_d)_{ij}}{\Delta_{ij}^{(d)}}\right| \quad \text{(up basis)}. \tag{9.32}$$

Using (9.31), this becomes an *overlap sum rule*

$$\theta_{ij}^{\mathrm{CKM}} \simeq \frac{1}{|\Delta_{ij}^{(u)}|}\left|\sum_m \lambda_m^{(u)}\, (a_m^{(u)})_i (a_m^{(u)})_j^*\right|, \tag{9.33}$$

making the hierarchy explicit: small misalignments of the dominant $a_3^{(u)}$ and suppressed subleading $\lambda_{1,2}^{(u)}$ generate the observed θ_{ij} pattern. The same holds with $u \leftrightarrow d$ in the other basis; consistency requires the two estimates to agree within higher orders.

Geometric CP measure (Jarlskog).

The Jarlskog invariant admits the exact identity

$$\mathcal{J} \prod_{i<j}(y_{u,j}^2 - y_{u,i}^2) \prod_{k<\ell}(y_{d,\ell}^2 - y_{d,k}^2) = \frac{1}{2i}\det\Big([H_u, H_d]\Big). \tag{9.34}$$

Substituting (9.31) yields a *purely geometric* expression: $\det([H_u, H_d])$ is a finite sum of triple products $\lambda_m^{(u)}\lambda_n^{(d)}\lambda_{m'}^{(u)}\,(a_m^{(u)}\cdot a_n^{(d)})\,(a_n^{(d)}\cdot a_{m'}^{(u)})\,(a_{m'}^{(u)}\cdot a_m^{(u)})$ whose imaginary part determines $\mathcal{J}$. Hence CP violation tracks oriented areas spanned by the overlap vectors in flavor space; it vanishes if all $a_m^{(u)}, a_n^{(d)}$ are coplanar or purely real.

Rank–one plus perturbations (Wolfenstein–like).

Let $H_{u,d} = \lambda_3^{(u,d)}\, \hat{a}_3^{(u,d)} \hat{a}_3^{(u,d)\dagger} + \delta H_{u,d}$, with $\|\delta H\| \ll \lambda_3$. Define the misalignment angle $\varphi := \arccos|\hat{a}_3^{(u)\dagger}\hat{a}_3^{(d)}|$. Then

$$\theta_{23}^{\mathrm{CKM}} \sim \varphi + \mathcal{O}\Big(\frac{\|\delta H\|}{\lambda_3}\Big), \qquad \theta_{13}, \theta_{12} = \mathcal{O}\Big(\frac{\|\delta H\|}{\lambda_3}\Big), \tag{9.35}$$

making the *overlap misalignment* the driver of V_{cb} and suppressed subleading overlaps responsible for V_{us}, V_{ub}.

Leptons: Dirac vs Majorana neutrinos.

For Dirac neutrinos one repeats (9.29)–(9.32) with $H_\nu = \widehat{Y}_\nu \widehat{Y}_\nu^\dagger$. For Majorana masses (Type-I seesaw)

$$M_\nu \;=\; -\frac{v_H^2}{2}\,\widehat{Y}_\nu\, M_R^{-1}\,\widehat{Y}_\nu^\top,$$

with M_R the RH-neutrino Majorana matrix (itself overlap–generated on $\mathcal{F}$ if desired). Let $U_\nu^\top M_\nu U_\nu = \mathrm{diag}(m_1, m_2, m_3)$. Large PMNS angles arise naturally if two (or more) overlap modes in M_ν are quasi-degenerate and geometrically misaligned with those of H_e.

Solar–reactor sum rule (generic form).

Write $U_{\mathrm{PMNS}} = U_e^\dagger U_\nu$. If U_e is a small rotation, $U_e \simeq R_{12}(\theta_{12}^e)R_{23}(\theta_{23}^e)R_{13}(\theta_{13}^e e^{-i\phi_e})$, and U_ν provides a leading structure (e.g. from two dominant overlap modes), then to first order

$$\theta_{12} \;\simeq\; \theta_{12}^\nu \;+\; \theta_{13}\cos\delta \;-\; \theta_{12}^e\cos\phi_e, \qquad \theta_{13} \;\simeq\; \left|\theta_{13}^\nu - \theta_{13}^e e^{i\phi_e}\right|, \tag{9.36}$$

where $(\theta_{12}^\nu, \theta_{13}^\nu)$ are determined by the eigenvectors of M_ν (hence by the overlap geometry in M_ν). Equation (9.36) is a *geometric* lepton sum rule: charged–lepton overlap rotations correct a neutrino–sector pattern (TBM, BM, GR, etc.) produced by $\mathcal{F}$.

Overlap control of PMNS phases.

Define the leptonic Jarlskog $\mathcal{J}_\ell = \Im[U_{e1}U_{\mu 2}U_{e2}^* U_{\mu 1}^*]$. In the Majorana case one has the exact relation

$$\mathcal{J}_\ell \prod_{i<j}(m_j^2 - m_i^2) \prod_{\alpha<\beta}(h_\beta^2 - h_\alpha^2) \;=\; \frac{1}{2i}\,\det\Big([M_\nu M_\nu^\dagger,\, H_e]\Big), \tag{9.37}$$

with (h_α) the charged-lepton Yukawa singular values. The RHS is again a finite sum of oriented overlap triple products from $\mathcal{F}$.

Texture–mixing correlations (predictive handles).

In an overlap basis where one mode dominates both sectors,

$$\widehat{Y}_u \approx y_t\, e_3 e_3^\top + \varepsilon_u\, X_u, \quad \widehat{Y}_d \approx y_b\, e_3 e_3^\top + \varepsilon_d\, X_d,$$

one finds, schematically,

$$\frac{y_c}{y_t} \sim \varepsilon_u, \quad \frac{y_s}{y_b} \sim \varepsilon_d, \quad \theta_{23}^{\mathrm{CKM}} \sim |\varepsilon_d\,(X_d)_{23} - \varepsilon_u\,(X_u)_{23}|, \quad \theta_{13}^{\mathrm{CKM}} \sim \left|\varepsilon_d\,(X_d)_{13} - \varepsilon_u\,(X_u)_{13}\right|, \tag{9.38}$$

tying mass hierarchies to mixing entries. Analogous relations hold in the lepton sector with X_e, X_ν and with M_ν replacing H_ν for Majorana masses.

Acceptance and deliverables (G1–G2, Y1–Y2).

- **G1/G2**: build $H_{u,d,e}$ (and M_ν) from profiles/modes on $\mathcal{F}$; audit positivity and unit consistency (Gate S).
- **Y1**: extract V_{CKM} and U_{PMNS}, report overlap sum rules (9.33), (9.34), (9.37); quantify geometric misalignment φ.
- **Y2**: propagate to low scales via RG (matrix running for $U(1)$ mixing), with uncertainties split into *geometry* (profiles, w), *thresholds/matching*, and *running*; provide DI-compatible error budgets as in §9.2.

Part V

Predictive Closure and Kill Signatures

Chapter 10

Compact Master Predictions

Prelude. This chapter compacts the geometric ingredients of Part IV into *predictive envelopes* for gauge couplings and the weak angle. At a reference scale μ_{ref} (declared on the DoV), the Lie–metric on the fiber fixes a positive gauge–kinetic kernel G and hence the bare inverse couplings. We then propagate these to M_Z with renormalization-group (RG) flow that includes (i) matrix running for abelian kinetic mixing, (ii) heavy-mode thresholds at DoV boundaries, and (iii) finite matching terms from CP embeddings. Uncertainties are carried throughout and split into *geometry* vs *running* vs *threshold/matching* components.

10.1 Gauge-coupling Ratios and $\sin^2\theta_W$ at a Reference Scale; RG to M_Z

Let $G = U(1)^m \times \prod_\alpha G_\alpha$ with Lie algebra basis $\{Y_I\}_{I=1}^m$ for the abelian sector and $\{T_a\}$ for simple factors. The positive gauge–kinetic kernel at μ_{ref} is

$$\mathsf{G}_{AB}(\mu_{\mathrm{ref}}) = \mathcal{J}\,\langle T_A, T_B\rangle_{\mathcal{F}}, \qquad \mathsf{G}(\mu_{\mathrm{ref}}) \succ 0, \tag{10.1}$$

with $\langle\cdot,\cdot\rangle_{\mathcal{F}}$ the fiber Lie–metric (§9.1.1) and $\mathcal{J} > 0$ a geometric scale (S0).

10.1.1 Reference-scale envelopes

Inverse couplings from geometry.

For a simple factor G_α choose a unit generator T_α (in the metric (10.1)). The inverse couplings at μ_{ref} are

$$\alpha_\alpha^{-1}(\mu_{\mathrm{ref}}) \;=\; 4\pi\,\mathcal{J}\,\langle T_\alpha, T_\alpha\rangle_{\mathcal{F}}, \qquad \mathsf{G}_{U(1)}(\mu_{\mathrm{ref}}) \;=\; \Big[\mathcal{J}\,\langle Y_I, Y_J\rangle_{\mathcal{F}}\Big]_{I,J=1}^{m}. \tag{10.2}$$

Let $q_Y \in \mathbb{R}^m$ select the hypercharge direction after any abelian mixing (§9.2). Then

$$u_{\rm ref} \;:=\; \frac{1}{g_Y^2(\mu_{\rm ref})} \;=\; q_Y^\top\, \mathsf{G}_{U(1)}^{-1}(\mu_{\rm ref})\, q_Y, \qquad v_{\rm ref} \;:=\; \frac{1}{g_2^2(\mu_{\rm ref})} \;=\; \mathcal{J}\,\langle T^L, T^L\rangle_{\mathcal{F}}. \tag{10.3}$$

The *geometric* weak angle at $\mu_{\rm ref}$ is therefore

$$\boxed{\sin^2\theta_W(\mu_{\rm ref}) \;=\; \frac{v_{\rm ref}}{u_{\rm ref}+v_{\rm ref}}, \qquad \alpha_3^{-1}(\mu_{\rm ref}) \;=\; 4\pi\,\mathcal{J}\,\langle T^C, T^C\rangle_{\mathcal{F}}.} \tag{10.4}$$

Envelope set and positivity gates.

Collect the geometric inputs as $\Theta_{\rm geom} = \Big(\mathsf{G}_{U(1)}(\mu_{\rm ref}), v_{\rm ref}, \alpha_3^{-1}(\mu_{\rm ref})\Big)$ with covariance $\mathrm{Cov}(\Theta_{\rm geom})$ obtained from the fiber model. Define the *reference-scale envelope* as the confidence ellipsoid

$$\mathcal{E}_{\rm ref}(p) \;:=\; \Big\{\Theta_{\rm geom} :\; (\Theta_{\rm geom}-\bar\Theta)^\top\,\Sigma^{-1}\,(\Theta_{\rm geom}-\bar\Theta) \;\le\; \chi^2_{p,k}\Big\}, \tag{10.5}$$

intersected with the positivity constraints $\mathsf{G}_{U(1)} \succ 0,\ v_{\rm ref} > 0$ (Gate S). Ratios such as α_3/α_2 and $\sin^2\theta_W$ map continuously from $\mathcal{E}_{\rm ref}(p)$ via (10.3)–(10.4).

Linearized uncertainties at $\mu_{\rm ref}$.

Let $s^2 := \sin^2\theta_W = v/(u+v)$ with $u = u_{\rm ref}$ and $v = v_{\rm ref}$. Then

$$\partial_u s^2 = -\frac{v}{(u+v)^2}, \qquad \partial_v s^2 = \frac{u}{(u+v)^2}, \qquad \delta u = -\mathrm{Tr}\Big(\mathsf{W}\,\delta\mathsf{G}_{U(1)}\Big),\ \ \mathsf{W} := \mathsf{G}_{U(1)}^{-1} q_Y q_Y^\top \mathsf{G}_{U(1)}^{-1}, \tag{10.6}$$

which yields $\mathrm{Var}[s^2(\mu_{\rm ref})]$ exactly as in (9.16). Report central values and 1σ bars; enforce $\sin^2\theta_W \in (0,1)$ automatically by construction.

10.1.2 RG maps and uncertainties

Affine RG map to M_Z.

Let $t = \ln(\mu/\mu_{\rm ref})$. For the SU(2) and SU(3) factors,

$$\alpha_a^{-1}(\mu) \;=\; \alpha_a^{-1}(\mu_{\rm ref}) \;-\; \frac{b_a}{2\pi}\,t \;-\; \sum_{t_k} \frac{\Delta b_a^{(k)}}{2\pi}\,\ln\frac{\mu}{M_k} \;-\; \delta_a^{\rm fin}, \qquad a \in \{2,3\}. \tag{10.7}$$

For the abelian block (kinetic mixing) one evolves the *inverse kernel*

$$\mathsf{G}_{U(1)}^{-1}(\mu) \;=\; \mathsf{G}_{U(1)}^{-1}(\mu_{\rm ref}) \;-\; \frac{\mathbf{B}_{U(1)}}{2\pi}\,t \;-\; \sum_{t_k} \frac{\Delta\mathbf{B}_{U(1)}^{(k)}}{2\pi}\,\ln\frac{\mu}{M_k} \;-\; \Delta^{\rm fin}, \tag{10.8}$$

and then project to hypercharge by $u(\mu) = q_Y^\top \mathsf{G}_{U(1)}^{-1}(\mu) q_Y$. The predicted weak angle at M_Z is the exact identity

$$\boxed{\sin^2\theta_W(M_Z) \;=\; \frac{v(M_Z)}{u(M_Z)+v(M_Z)}, \qquad v(M_Z) = \alpha_2^{-1}(M_Z)/(4\pi).} \tag{10.9}$$

Covariance propagation.

Stack the *reference* variables in a vector

$$X_{\text{ref}} = \Big(\alpha_2^{-1},\ \alpha_3^{-1},\ \text{vec}[\mathsf{G}_{U(1)}^{-1}]\Big)_{\mu_{\text{ref}}},$$

with covariance Σ_{ref}. The RG map $X(M_Z) = A\,X_{\text{ref}} + b + \sum_k C_k \ln(M_Z/M_k)$ is affine; hence

$$\Sigma_X(M_Z) \;=\; A\,\Sigma_{\text{ref}}\,A^\top \;+\; \Sigma_{\text{run}} \;+\; \Sigma_{\text{thr}}, \tag{10.10}$$

where Σ_{run} encodes scheme/loop truncation and Σ_{thr} collects threshold/matching covariances from $\{\Delta b, \Delta\mathbf{B}, \delta^{\text{fin}}, \Delta^{\text{fin}}\}$. Finally,

$$\text{Var}\Big[\sin^2\theta_W(M_Z)\Big] \;=\; (\nabla s^2)^\top\,\text{Cov}\Big(u(M_Z), v(M_Z)\Big)\,(\nabla s^2),$$

with ∇s^2 as in (10.6) evaluated at M_Z.

Positivity and gates along the flow.

Enforce for all $\mu \in [M_Z, \mu_{\text{ref}}]$:

$$\mathsf{G}_{U(1)}(\mu) \succ 0, \qquad \alpha_2^{-1}(\mu) > 0,\ \alpha_3^{-1}(\mu) > 0, \tag{10.11}$$

(Gate S), and implement ST-consistent matching at each M_k (Gate TQ). Locality gates G restrict thresholds to declared DoV boundaries (appearing/disappearing modes).

Deliverables (for this section).

(i) Reference-scale envelope $\mathcal{E}_{\text{ref}}(p)$ for $\{\alpha_2^{-1}, \alpha_3^{-1}, \mathsf{G}_{U(1)}\}$ and induced ranges for $\sin^2\theta_W(\mu_{\text{ref}})$; (ii) RG map to M_Z with $\Sigma_X(M_Z)$ and the final $\sin^2\theta_W(M_Z)$ prediction with a split uncertainty budget (geometry / running / thresholds); (iii) positivity and gate-compliance checks (10.11).

10.2 Precision Baselines: (S, T, U) and EDM Floors

Prelude. To turn the MSP geometry into falsifiable statements we need (i) a compact, scheme–transparent definition of the *oblique* parameters (S, T, U) that capture universal gauge–boson self–energy shifts on (M, g_{eff}), and (ii) *irreducible electric dipole moment (EDM) floors*: the minimal, model–independent CP–odd signals implied by the over-

lap–driven Yukawa phases (§9.3) and by gauge–invariant operators on the DoV. Both sets serve as *kill signatures*: excursions beyond the acceptance windows (Chapter 11) falsify the construction.

10.2.1 Oblique parameters

Vacuum polarizations and conventions.

Let $\Pi^{\mu\nu}_{XY}(q)$ be the 1PI two–point function of gauge currents $X, Y \in \{W^1, W^2, W^3, B\}$ on (M, g_{eff}) with $\Pi^{\mu\nu}_{XY}(q) = (q^\mu q^\nu - q^2 g^{\mu\nu}_{\text{eff}})\,\Pi_{XY}(q^2)$. Derivatives are denoted $\Pi'_{XY}(0) := \mathrm{d}\Pi_{XY}(q^2)/\mathrm{d}q^2|_{q^2=0}$. Write $\alpha := e^2/(4\pi)$, $s_W := \sin\theta_W$, $c_W := \cos\theta_W$.

Peskin–Takeuchi definition (gauge basis).

On the DoVwe adopt

$$S \;:=\; 16\pi\,\Pi'_{3B}(0), \qquad T \;:=\; \frac{4\pi}{s_W^2 c_W^2 m_Z^2}\Big(\Pi_{11}(0)-\Pi_{33}(0)\Big), \qquad U \;:=\; 16\pi\,\Big(\Pi'_{11}(0)-\Pi'_{33}(0)\Big). \tag{10.12}$$

Equivalently, in the mass basis,

$$\alpha\,T \;=\; \frac{\Pi_{WW}(0)}{m_W^2} - \frac{\Pi_{ZZ}(0)}{m_Z^2}, \qquad \alpha\,S \;=\; 4 s_W c_W\,\Pi'_{Z\gamma}(0), \qquad \alpha\,(S+U) \;=\; 4 s_W^2\,\Pi'_{WW}(0). \tag{10.13}$$

These definitions are scheme–stable to the order used and respect Ward/ST identities (Gate TQ).

Geometry → oblique structures.

MSP fixes the gauge–kinetic kernel at μ_{ref}, $\mathscr{L}_{\text{YM}} = -\frac{1}{4}\mathsf{G}_{AB}F^A_{\mu\nu}F^{B\,\mu\nu}$, with $\mathsf{G} \succ 0$. Small geometric deformations $\delta\mathsf{G}$ and mixing with heavy modes across DoVboundaries (thresholds) yield linear responses

$$\delta S \;=\; \boldsymbol{s}\cdot\mathrm{vec}(\delta\mathsf{G}) \;+\; \sum_t \boldsymbol{s}_t\,\ln\frac{\mu}{M_t}, \qquad \delta T \;=\; \boldsymbol{t}\cdot\mathrm{vec}(\delta\mathsf{G}) \;+\; \sum_t \boldsymbol{t}_t\,\ln\frac{\mu}{M_t}, \tag{10.14}$$

(and similarly for U), where the coefficients follow from matching the polarizations induced by $\delta\mathsf{G}$ and by integrating out the heavy pseudomodes (local CP embeddings, Chapter 5). Equation (10.14) is the model's *master map* from geometry to (S, T, U).

Low–energy parameterization (EFT view, schematic).

In a gauge–invariant basis one may summarize leading corrections by

$$\Delta\mathscr{L} \;\supset\; \frac{c_{WB}}{\Lambda^2}(H^\dagger\tau^a H)\,W^a_{\mu\nu}B^{\mu\nu} \;+\; \frac{c_{HD}}{\Lambda^2}\,(H^\dagger D_\mu H)(H^\dagger D^\mu H) + \cdots, \tag{10.15}$$

which induce, at tree level,

$$S \;\propto\; 16\pi\,\frac{v^2}{\Lambda^2}\,c_{WB}, \qquad T \;\propto\; -\,\frac{4\pi}{\alpha}\,\frac{v^2}{\Lambda^2}\,c_{HD}, \tag{10.16}$$

up to known normalization factors. (*We keep the mapping symbolic since normalization conventions vary; the MSP pipeline computes the precise coefficients from* G *and threshold data.*)

Uncertainties and acceptance.

Define the covariance at μ:

$$\Sigma_{STU}(\mu) \;=\; A\,\Sigma_{\text{geom}}\,A^{\top} \;+\; \Sigma_{\text{thr}} \;+\; \Sigma_{\text{run}},$$

where A collects the linear map in (10.14). *Acceptance* requires reporting central (S, T, U), their covariance, and positivity of the propagators (no tachyon/ghost in gauge sector) throughout the DoV(Gate S). Compatibility with (S, T, U) reference ellipses at M_Z becomes a kill test in Chapter 11.

10.2.2 Irreducible EDM floors

CP–odd effective operators (basis).

At hadronic/leptonic scales, the leading CP–odd terms are

$$\mathscr{L}_{\text{EDM}} = -\,\frac{i}{2}\sum_f d_f\,\bar{f}\,\sigma^{\mu\nu}\gamma_5 f\,F_{\mu\nu} - \frac{i}{2}\sum_q \tilde{d}_q\,\bar{q}\,\sigma^{\mu\nu}\gamma_5 T^a q\,G^a_{\mu\nu} + \frac{w}{6}\,f^{abc}G^a_{\mu\rho}G^{b\,\rho}{}_{\nu}\,\widetilde{G}^{c\,\mu\nu} \;+\; \cdots\,, \tag{10.17}$$

with d_f (leptonic/quark EDMs), $\tilde{d}_q$ (quark chromo–EDMs), and w (Weinberg three–gluon operator), and $\widetilde{G}^{\mu\nu} = \frac{1}{2}\epsilon^{\mu\nu\rho\sigma}G_{\rho\sigma}$. Wilson coefficients *at the matching scale* are generated by heavy loops sourced by CP phases in Yukawas/gauge–Higgs couplings; RG running mixes the operators down to hadronic scales.

MSP sources of irreducible CP.

Within MSP, overlap–driven Yukawas (§9.3) generically carry rephasing–invariant phases (Jarlskog–type), hence—absent an additional symmetry— there exist *irreducible* two–loop Barr–Zee–type contributions

$$d_\ell \;\sim\; \frac{e\,\alpha}{8\pi^3}\,\frac{m_\ell}{m_H^2}\sum_i Q_i^2\,\Im(\kappa_i)\,F\!\left(\frac{m_i^2}{m_H^2}\right), \qquad d_q,\;\tilde{d}_q \text{ analogous},$$

where i runs over heavy charged states (e.g. t, W), κ_i encode the effective CP–odd Higgs–fermion/vector couplings determined by the overlap phases, and F are loop functions. These terms define the *floor*: they cannot be dialed to zero without removing the geometric

CP invariants.

RG transport and master formulae.

Stack the short–distance Wilson vector $C(\mu_h) = (d_f, \tilde{d}_q, w, \ldots)_{\mu_h}$ at a high scale μ_h set by the heaviest loop. Evolve to hadronic/leptonic scales by

$$C(\mu_\ell) \;=\; U(\mu_\ell, \mu_h)\, C(\mu_h), \qquad U \;=\; \mathcal{T}\exp\!\left(\int_{\mu_\ell}^{\mu_h} \frac{\mathrm{d}\mu}{\mu}\, \gamma^T(\mu)\right), \tag{10.18}$$

with γ the anomalous–dimension matrix (includes mixing among $\{d, \tilde{d}, w\}$). Observable EDMs follow from hadronic/nuclear/atomic matrix elements, e.g.

$$\begin{aligned} d_n &= a_u d_u + a_d d_d + a_s d_s + b_u \tilde{d}_u + b_d \tilde{d}_d + b_s \tilde{d}_s + c\, w + \cdots\,, \\ d_{\mathrm{Hg}},\ d_{\mathrm{Ra}} &:\ \text{linear combinations of } \{d_q, \tilde{d}_q, w\} \text{ through Schiff moments,} \end{aligned} \tag{10.19}$$

with (a_i, b_i, c) obtained from QCD sum rules or lattice inputs. Leptonic d_e is perturbative with small atomic/nuclear enhancements (paramagnetic systems).

Geometric lower bounds (floors).

Let $\mathcal{J}_{\mathrm{Q}}$, $\mathcal{J}_{\mathrm{L}}$ be the quark/lepton Jarlskog–type invariants computed from the overlap Yukawas at μ_{ref}. Then, for any choice of geometric weights w that reproduces observed masses and mixings,

$$|d_e|_{\mathrm{floor}} \;\gtrsim\; K_e \left|\mathcal{J}_{\mathrm{L}}\right|, \qquad |d_n|_{\mathrm{floor}} \;\gtrsim\; K_n \left|\mathcal{J}_{\mathrm{Q}}\right|, \tag{10.20}$$

with $K_{e,n}$ calculable positive functionals of loop factors, heavy masses and RG/hadronic dressing (pipeline outputs). The inequalities become equalities if no accidental cancellations occur among multiple Barr–Zee topologies. Any experimental bound *below* (10.20) is a *kill* for the overlap CP structure on the DoV.

Uncertainties and acceptance.

Split errors as

$$\Delta d \;=\; \sqrt{(\Delta d)^2_{\mathrm{geom}} + (\Delta d)^2_{\mathrm{RG}} + (\Delta d)^2_{\mathrm{had}}},$$

where *geom* propagates the profile/mode uncertainties (Yukawa phases and magnitudes), *RG* from scale variation and truncation in U, and *had* from matrix–element determinations. *Acceptance* requires reporting $(d_e, d_n, \ldots)$ with the split error budget and verifying that CP–odd operators respect gauge invariance/positivity and locality on the DoV(Gates $\mathsf{TQ}, \mathsf{S}, \mathsf{G}$). Floors (10.20) must be published as numerical envelopes alongside (S, T, U).

Artifact 10.1: AR: Precision baselines

Artifacts: (i) a polarization module computing Π_{XY} and (S, T, U) from $\delta\mathsf{G}$ and thresholds; (ii) an EDM builder generating $C(\mu_h)$ from overlap Yukawas, transporting via (10.18), and assembling observables via (10.19); (iii) uncertainty budget split (geom/RG/had) and gate–compliance checks; (iv) regression tests against known EFT limits.

10.3 Neutrino NSI/CPT-odd Signatures and Appearance Channels

Scope. On the declared DoVand backgrounds (M, g_{eff}), we develop an *operational* Hamiltonian for flavor evolution that augments the standard three–flavor dynamics by (i) non-standard interactions (NSI) in matter and (ii) CPT-odd contributions compatible with locality/CP gates. We extract *matrix patterns* for the effective NSI couplings from fiber geometry and overlap arguments (MSP), and identify *waveform* features in appearance probabilities vs. baseline L and energy E that furnish kill signatures and DI-friendly likelihoods.

10.3.1 NSI matrix patterns

Effective Hamiltonian on (M, g_{eff}).

Write the flavor–basis Schrödinger equation

$$i\,\frac{\mathrm{d}}{\mathrm{d}x}\,\nu_\alpha(x) \;=\; H_{\text{eff}}(x,E)\,\nu_\alpha(x), \qquad H_{\text{eff}} \;=\; H_{\text{vac}}(E) + V_{\text{mat}}(x) + V_{\text{NSI}}(x) + H_{\text{CPT}}(E), \tag{10.21}$$

with

$$H_{\text{vac}}(E) = \frac{1}{2E}\,U\,\mathrm{diag}(0, \Delta m^2_{21}, \Delta m^2_{31})\,U^\dagger, \tag{10.22}$$

$$V_{\text{mat}}(x) = \mathrm{diag}\Big(V_{CC}(x), 0, 0\Big), \quad V_{CC}(x) := \sqrt{2}\,G_F\,N_e(x), \tag{10.23}$$

$$V_{\text{NSI}}(x) = V_{CC}(x)\,\varepsilon, \qquad \varepsilon = \varepsilon^\dagger \in \mathbb{C}^{3\times 3}, \quad (\varepsilon_{\alpha\alpha} \in \mathbb{R}), \tag{10.24}$$

$$H_{\text{CPT}}(E) = a - E\,c, \qquad a = a^\dagger, \quad c = c^\dagger \text{ (dimensionful Hermitian } 3\times 3), \tag{10.25}$$

where U is the PMNS matrix (PDG phase convention). The NSI matrix ε collects neutral-current effective interactions normalized to V_{CC} (so that $\varepsilon \equiv 0$ reproduces the standard MSW term). CPT–odd contributions are parameterized by a (dimension of energy; flips sign under $\nu \leftrightarrow \bar{\nu}$) and c (dimensionless; modifies dispersion), both constrained by Gates TQ and S (positivity, no superluminal signalling; §4.2).

Geometric/overlap origin and patterns.

MSP overlap arguments (§9.3) imply that forward–scattering amplitudes in matter inherit *low-rank textures*. A generic, symmetry–tolerant ansatz on the DoVis

$$\varepsilon \;=\; \sum_{m=1}^{r} \lambda_m\, \mathbf{u}_m\, \mathbf{u}_m^\dagger, \qquad r \le 3, \quad \lambda_m \in \mathbb{R}, \quad \mathbf{u}_m \in \mathbb{C}^3, \quad \|\mathbf{u}_m\| = 1, \tag{10.26}$$

with phases in $\mathbf{u}_m$ traceable to fiber profiles. Rank–one dominance ($r = 1$) yields *alignment patterns* $\varepsilon \propto \mathbf{u}\mathbf{u}^\dagger$, while residual fiber symmetries enforce *texture zeros* (e.g. $(\varepsilon)_{e\tau} = 0$) via Schur selection as in (9.26).

Rephasing invariants and priors.

The physically relevant NSI phases are rephasing invariants constructed from ε and U:

$$\phi_{e\mu} := \arg(\varepsilon_{e\mu} U_{e3}^* U_{\mu 3}), \quad \phi_{e\tau} := \arg(\varepsilon_{e\tau} U_{e3}^* U_{\tau 3}), \quad \ldots$$

We adopt priors on $(\lambda_m, \mathbf{u}_m)$ from the fiber metric on $\mathcal{F}$ and the same weight w used in Yukawa overlaps; these priors propagate to Gaussian/elliptic envelopes for the entries of ε.

Hermiticity/positivity gates.

Gate S demands H_{eff} self-adjoint and the evolution unitary on the DoV. This yields the constraints

$$\varepsilon = \varepsilon^\dagger, \qquad a = a^\dagger, \qquad c = c^\dagger, \qquad \|c\| < 1 \ \text{(no superluminal group velocities on } (M, g_{\text{eff}})).$$

Gate TQ enforces CP consistency: all effective terms arise from local, gauge-invariant operators after integrating out heavy modes (cf. Chapter 5).

First–order appearance shifts (constant density, leading order).

For constant N_e and small parameters $\alpha := \Delta m_{21}^2/\Delta m_{31}^2$, s_{13}, $\varepsilon_{\alpha\beta}$, a/E, and c, the $\nu_\mu \to \nu_e$ probability at baseline L is

$$P_{\mu e} = P_{\mu e}^{(0)} \;+\; \Delta P_{\mu e}^{\text{NSI}} \;+\; \Delta P_{\mu e}^{\text{CPT}} \;+\; \mathcal{O}(\text{quadratic}), \tag{10.27}$$

$$P_{\mu e}^{(0)} = 4 s_{13}^2 s_{23}^2 \frac{\sin^2\!\big((1-\hat{A})\Delta\big)}{(1-\hat{A})^2} + \alpha^2 \sin^2 2\theta_{12}\, c_{23}^2 \frac{\sin^2(\hat{A}\Delta)}{\hat{A}^2} + \alpha s_{13} \sin 2\theta_{12} \sin 2\theta_{23} \frac{\sin(\hat{A}\Delta)}{\hat{A}} \frac{\sin\!\big((1-\hat{A}\big.}{(1-\hat{A})} \tag{10.28}$$

$$\Delta P_{\mu e}^{\text{NSI}} = 4\hat{A}\, s_{13} \frac{\sin(\hat{A}\Delta)}{\hat{A}} \frac{\sin\!\big((1-\hat{A})\Delta\big)}{(1-\hat{A})} \Big[c_{23}\, |\varepsilon_{e\mu}| \cos(\Delta + \delta - \phi_{e\mu}) - s_{23}\, |\varepsilon_{e\tau}| \cos(\Delta + \delta - \phi_{e\tau}) \Big]$$
$$+\; 2\hat{A}\, \sin 2\theta_{23}\, |\varepsilon_{\mu\tau}|\, \sin^2\!\Delta\, \cos\phi_{\mu\tau} \;+\; \hat{A}\, (\varepsilon_{ee} - \varepsilon_{\mu\mu}) \frac{\partial P_{\mu e}^{(0)}}{\partial \hat{A}}, \tag{10.29}$$

$$\Delta P_{\mu e}^{\text{CPT}} = 2L \Big[\Re(a_{e\mu})\, \mathcal{F}_1(E, L) - \Im(a_{e\mu})\, \mathcal{F}_2(E, L) \Big] \;+\; \mathcal{O}(c), \tag{10.30}$$

where

$$\Delta := \frac{\Delta m_{31}^2 L}{4E}, \qquad \hat{A} := \frac{2EV_{CC}}{\Delta m_{31}^2}, \qquad \mathcal{F}_{1,2}(E,L) \text{ are bounded oscillatory kernels (known functions).}$$

For antineutrinos: $\delta \to -\delta$, $\hat{A} \to -\hat{A}$, $\varepsilon \to \varepsilon^*$, $a \to -a$, $c \to c$.

Degeneracies and geometric lifts.

The *LMA-dark*–like transformation $\theta_{12} \to \frac{\pi}{2} - \theta_{12}$, $\delta \to \pi - \delta$, $(\varepsilon_{ee} - \varepsilon_{\mu\mu}) \to (\varepsilon_{ee} - \varepsilon_{\mu\mu}) - 2$ leaves leading terms nearly invariant in constant density. Geometric priors (10.26)—particularly rank–one alignment—break this degeneracy by correlating $(\varepsilon_{e\mu}, \varepsilon_{e\tau})$ with diagonal entries and restricting allowed phases.

Robustness to colored/cone slacks.

Colored dynamics and finite-speed leakage from §7.3 induce additive DI slacks in event rates:

$$|P_{\alpha\beta}^{\text{obs}} - P_{\alpha\beta}^{\text{th}}| \;\le\; \Delta_{\text{stat}} \;+\; \underbrace{2\,\zeta_{\text{col}}}_{\text{colored}} \;+\; \underbrace{\epsilon_{\text{cone}}}_{\text{cone}},$$

which we include in envelopes and likelihoods.

10.3.2 Waveform shapes & baselines

Energy–baseline *waveforms*.

Define the appearance surface $P_{\mu e}(L,E)$ and its CP–conjugate $\bar{P}_{\mu e}(L,E)$. NSI and CPT-odd terms induce characteristic distortions:

- **NSI phase–shifted fringes:** off-diagonal $|\varepsilon_{e\mu}|, |\varepsilon_{e\tau}|$ produce phase shifts $\Delta \to \Delta + \varphi_{\text{NSI}}$ in the leading fringe via (10.29), with opposite sign between ν and $\bar{\nu}$ through $\hat{A}$.
- **Diagonal tilt:** $(\varepsilon_{ee} - \varepsilon_{\mu\mu})$ effectively rescales $\hat{A}$, tilting the nodal lines in (L,E).
- **CPT-odd odd–under–CP ridge:** $a_{e\mu}$ yields a linear-in-L odd contribution (10.30) that *reverses sign* between ν and $\bar{\nu}$, thus creating a ridge/valley asymmetry not reproducible by δ alone.

Magic baselines (phase decoupling).

For constant density, the δ-dependence of $P_{\mu e}^{(0)}$ vanishes when $\sin(\hat{A}\Delta) = 0$, i.e.

$$\int_0^{L_{\text{mag}}} V_{CC}(x)\,\mathrm{d}x \;=\; n\pi\,\frac{\Delta m_{31}^2}{2E}, \qquad n \in \mathbb{Z}. \tag{10.31}$$

At such L_{mag}, NSI off-diagonals still enter through (10.29), allowing a clean $\{\varepsilon\}$ probe with reduced δ degeneracy. For variable density, replace the integral by the path-ordered exponential along the geodesic in (M, g_{eff}).

Near–far and dual–baseline strategies.

Choose (L_1, E_1) close to the first oscillation maximum ($\Delta \sim \pi/2$) and (L_2, E_2) near a magic line (10.31). The pair $\{P_{\mu e}(L_1, E_1),\ P_{\mu e}(L_2, E_2)\}$ breaks the $\{\delta, \varepsilon_{e\mu}, \varepsilon_{e\tau}\}$ degeneracy; adding $\bar{P}_{\mu e}$ constrains $a_{e\mu}$ via the CP–odd ridge.

Two-flavor CPT-odd diagnostic (analytic).

In an effective two-flavor slice with mixing θ, splitting Δm^2, constant a along the flavor axis, one finds

$$P_{\alpha\beta} - \bar{P}_{\alpha\beta} \;\simeq\; 4\,\frac{aL}{2E}\,\sin 2\theta\,\sin^2\Delta \;+\; \mathcal{O}(a^2), \qquad \Delta = \frac{\Delta m^2 L}{4E}, \tag{10.32}$$

a clean odd-under-CP diagnostic insensitive to δ at leading order.

Likelihoods and envelopes.

Define the parameter vector $\vartheta = (\Delta m^2_{ij}, \theta_{ij}, \delta, \varepsilon, a, c)$. Given binned spectra $\mathcal{D}$ and transfer matrices (flux→rates), the end-to-end likelihood is

$$\mathcal{L}(\mathcal{D} \mid \vartheta) = \prod_b \mathrm{Pois}\Big(n_b^{\mathrm{obs}}\,;\, n_b^{\mathrm{th}}(\vartheta)\Big) \times \mathcal{N}(\mathrm{sys} \mid 0, \Sigma_{\mathrm{sys}}) \times \mathbf{1}_{\mathrm{gates}}(\vartheta),$$

with $\mathbf{1}_{\mathrm{gates}}$ enforcing $\mathsf{S}, \mathsf{TQ}, \mathsf{G}$ and DI slacks. Marginalizing delivers envelope plots in $(\varepsilon_{\alpha\beta}, a_{\alpha\beta})$ and correlations with δ.

Deliverables.

Numeric envelopes for ε and a (with geometric priors), correlation matrices (ε, δ) and (a, δ), and end-to-end likelihoods that ingest (L, E)-binned counts, include colored/cone slacks, and output DI-certified credible regions for appearance/disappearance channels.

Chapter 11

Falsifiability and Acceptance Windows

Prelude. This chapter formalizes the *kill signatures* of the framework. On the declared DoV, acceptance requires two complementary families of tests: (i) *forward-limit* constraints from analyticity, unitarity, and crossing of $2\to 2$ amplitudes on (M, g_{eff}); and (ii) *cone-LMI* (Local Matrix Inequality) feasibility conditions ensuring hyperbolic evolution and microscopic causal support within the emergent cone of g_{eff}. Violations of either family falsify MSP on the given DoV(Gate S/G).

11.1 Positivity/Hyperbolicity Gates; Cone-LMI Feasibility Regions

Scope. We translate S-matrix principles (analyticity, unitarity, crossing) and hyperbolicity of the field equations into *computable* inequalities on effective couplings and kinetic kernels. The forward-limit bounds act on Wilson coefficients of higher-derivative operators; cone-LMIs constrain the principal symbol and characteristic cone of propagation. Together they define a feasibility region that must be nonempty on the DoV.

11.1.1 Forward-limit constraints

Kinematics on (M, g_{eff}).

Let p_i^μ be on-shell momenta of external one-particle states supported in pointer sectors (Gate S). With the emergent metric g_{eff}, define Mandelstam invariants

$$s := -(p_1+p_2)^2_{g_{\text{eff}}}, \quad t := -(p_1-p_3)^2_{g_{\text{eff}}}, \quad u := -(p_1-p_4)^2_{g_{\text{eff}}}, \qquad s+t+u = \sum_i m_i^2, \tag{11.1}$$

and consider elastic $2 \to 2$ amplitudes $\mathcal{M}_{IJ}(s,t)$ where I, J label incoming/outgoing flavor–helicity channels (projected with physical polarizations). The DoVis chosen so that branch points start at $s \geq s_{\rm th} > 0$ and the amplitude obeys polynomial boundedness in s.

Analyticity, unitarity, crossing (Gates TQ/S).

Hermitian analyticity and crossing hold in the forward limit $t \to 0$:

$$\mathcal{M}_{IJ}(s,0) \;=\; \mathcal{M}_{JI}(s,0), \qquad \Im\,\mathcal{M}_{II}(s,0) \;=\; s\,\sigma_I^{\rm tot}(s) \;\geq 0 \quad (s > s_{\rm th}), \tag{11.2}$$

and $\mathcal{M}_{IJ}(s,0)$ is analytic in the complex s-plane cut along $[s_{\rm th}, \infty)$ and $(-\infty, -s_{\rm th}]$.

Subtracted dispersion and positivity.

Assuming N subtractions (Froissart–Martin type bound on the DoV), fix $s_* < s_{\rm th}$ and write for each channel pair (I, J)

$$\partial_s^2 \mathcal{M}_{IJ}(s_*,0) \;=\; \frac{2}{\pi}\int_{s_{\rm th}}^{\infty} \mathrm{d}s' \, \frac{\Im\,\mathcal{M}_{IJ}(s',0)}{(s'-s_*)^3} \;+\; \frac{2}{\pi}\int_{s_{\rm th}}^{\infty} \mathrm{d}s' \, \frac{\Im\,\mathcal{M}_{IJ}(-s',0)}{(s'+s_*)^3}. \tag{11.3}$$

Using crossing to relate the left cut to the right cut and $\Im\,\mathcal{M} \succeq 0$ from (11.2), one gets for any complex test vector $v \in \mathbb{C}^{n_{\rm ch}}$

$$v^\dagger\, \partial_s^2 \mathcal{M}(s_*,0)\, v \;\geq\; 0, \qquad \text{i.e.} \quad \partial_s^2 \mathcal{M}(s_*,0) \;\succeq\; 0. \tag{11.4}$$

This *matrix* positivity is the central forward-limit constraint.

EFT projection and Wilson-coefficient LMIs.

Expand the low-energy amplitude at $t = 0$ (below the lowest inelastic threshold)

$$\mathcal{M}_{IJ}(s,0) \;=\; \mathcal{M}_{IJ}^{(2)}\, s \;+\; \mathcal{M}_{IJ}^{(4)}\, s^2 \;+\; \sum_{k\geq 3} \mathcal{C}_{IJ}^{(2k)}\, s^k \;+\; \cdots, \tag{11.5}$$

where $\mathcal{C}^{(2k)}$ are linear combinations of Wilson coefficients of local operators $\mathcal{O}^{(2k)}$ (dimension $2k+4$) after integrating out heavy pseudomodes (CP embeddings; Chapter 5). Then (11.4) implies the *EFT positivity LMIs*

$$\boxed{\mathcal{C}^{(6)} \;\succeq\; 0, \quad \mathcal{C}^{(8)} \;\succeq\; 0, \;\ldots \quad \text{in every elastic channel block and polarization sector.}} \tag{11.6}$$

Concretely, for a basis of external polarizations $\{\epsilon_r\}$, the quartic gauge/tensor structures project to Gram matrices built from $(\epsilon_r \cdot \epsilon_s)$ and $(\epsilon_r \cdot p_i)$, and (11.6) requires all such Gram-contracted coefficient matrices to be PSD.

Examples (schematic).

- *Scalar quartic derivatives:* $\Delta\mathscr{L} \supset c^{(8)}_{ijkl}(\partial\phi_i\cdot\partial\phi_j)(\partial\phi_k\cdot\partial\phi_\ell)$. Then $V^\top c^{(8)} V \geq 0$ for all vectors V with components $V_{ij} = \alpha_i\alpha_j$ (elastic forward kinematics), requiring $c^{(8)} \succeq 0$ on the symmetric tensor subspace.
- *Gauge quartics:* $\Delta\mathscr{L} \supset \sum_a \left(f_1^{(8)}(F^a_{\mu\nu}F^{a\,\mu\nu})^2 + f_2^{(8)}(F^a_{\mu\nu}\widetilde{F}^{a\,\mu\nu})^2 + \cdots\right)$. Forward helicity amplitudes enforce $f^{(8)}_{1,2} \geq 0$ and PSD constraints on mixed-color blocks (after projecting to physical, transverse polarizations on (M, g_{eff})).
- *Fermion contact terms:* for $(\bar{\psi}\gamma_\mu\psi)^2$-type operators, the spin-averaged forward amplitude yields PSD conditions on the flavor-current coefficient matrices.

Gravity/effective cone consistency.

If the EFT includes curvature couplings on (M, g_{eff}), the same dispersion argument applies to eikonalized matter scattering in weakly curved regions of the DoV. Violations (e.g. superluminal phase/group speeds for some polarization) signal either breakdown of the DoVor inconsistency of the sign pattern in $\mathcal{C}^{(2k)}$; these are *kill signatures* (Gate G/S).

Finite-t robustness (beyond forward).

For small spacelike $t < 0$, analyticity allows (11.3) at $t \neq 0$, leading to families of inequalities on mixed derivatives $\partial_s^m \partial_t^\ell \mathcal{M}(s_*, 0)$ that bound combinations of Wilson coefficients multiplying $(\partial\cdot\partial)^{m+\ell}$ operators. We record these as optional *strengthened* gates.

Acceptance checks and numerics.

Checklist 11.1: Forward-limit gate (positivity)

- Build channel–polarization blocks of $\mathcal{M}(s, 0)$ at low energy from the EFT on the DoV; extract $\{\mathcal{C}^{(2k)}\}$.
- Assemble the Gram-projected matrices and verify $\mathcal{C}^{(2k)} \succeq 0$ (LMIs) for all tested directions.
- Audit dispersion prerequisites: location of s_{th}, boundedness, and optical-theorem positivity (11.2).
- Flag any violation as a *kill* for the given DoV; otherwise report the minimal eigenvalues as *margins*.

Artifact 11.1: AR: Forward-limit positivity suite

Scripts to (i) generate helicity/polarization projectors on (M, g_{eff}), (ii) build $\mathcal{M}(s, 0)$ from the EFT to $\mathcal{O}(s^2)$ and beyond, (iii) assemble/solve the LMIs (11.6) (reporting smallest eigenvalues and offending directions), and (iv) regression tests on toy models saturating the bounds.

11.1.2 Cone LMIs

First–order symmetric–hyperbolic form.

On a local orthonormal frame of (M, g_{eff}) with $g_{\text{eff}} = \text{diag}(-1, 1, 1, 1)$, write the linearized equations for a field multiplet $u(x) \in \mathbb{R}^N$ in first–order form

$$A^0(x)\,\partial_t u + A^i(x)\,\partial_i u + B(x)\,u \;=\; J(x), \qquad i = 1, 2, 3, \tag{11.7}$$

with real $N \times N$ matrices A^μ. On the declared DoVwe treat A^μ as slowly varying (frozen–coefficient analysis).

Symmetrizer LMIs (strong hyperbolicity).

A sufficient (and standard) condition for strong/symmetric hyperbolicity is the existence of a *symmetrizer* $H(x) = H(x)^\top \succ 0$ such that

$$HA^0 \;\succeq\; m\,\mathbb{I}, \qquad HA^i \;=\; (HA^i)^\top \quad (i = 1, 2, 3), \tag{11.8}$$

for some $m > 0$. The constraints (11.8) are linear equalities plus the LMI $H \succ 0$ and can be solved by SDP on each slab of the DoV.

Gravity cone (dominance).

Let γ_{ij} be the spatial metric induced by g_{eff} in the chosen frame ($g_{\text{eff}} = -\mathrm{d}t^2 + \gamma_{ij}\mathrm{d}x^i\mathrm{d}x^j$). For any unit spatial direction $n \in \mathbb{R}^3$ with $\|n\|_\gamma^2 := \gamma_{ij}n^i n^j = 1$, define $A(n) := \sum_i A^i n_i$. Subluminal propagation w.r.t. the g_{eff} light cone is ensured by the *cone-dominance* block LMIs

$$\boxed{\begin{bmatrix} HA^0 & HA(n) \\ A(n)^\top H & HA^0 \end{bmatrix} \;\succeq\; 0 \quad \text{for all unit } n \text{ (in } \gamma\text{–norm).}} \tag{11.9}$$

Equivalently, (11.9) is the operator inequality $A(n)^\top H(A^0)^{-1}H^{-1}A(n) \;\preceq\; A^0$ (in the H–metric), implying $\rho\Big((A^0)^{-1/2}A(n)(A^0)^{-1/2}\Big) \le 1$, hence the characteristic speeds do not exceed the g_{eff} light speed. Together, (11.8)–(11.9) guarantee that the matter characteristic cone is contained in (or equal to) the gravity cone.

Eikonal/principal–symbol view (second–order systems).

For a second–order kinetic operator with principal symbol $P(\xi) = K^{\mu\nu}\xi_\mu\xi_\nu$ (acting on $\mathbb{R}^N$), reduce to first order or enforce the following PSD conditions directly:

$$\text{for all null } \xi\ (g_{\text{eff}}^{\mu\nu}\xi_\mu\xi_\nu = 0): \quad v^\top P(\xi)\,v \;\ge\; 0\ \forall v \in \mathbb{R}^N, \qquad \text{and for timelike } t:\ v^\top P(t)\,v \;\ge\; m\,\|v\|^2. \tag{11.10}$$

On a null cone, (11.10) prevents superluminal phase speeds; on timelike t, it yields coercivity (well–posed Cauchy problem).

S-procedure / sampled directions.

The quantified constraint in (11.9) over the unit sphere can be implemented (i) by dense sampling $\{n_k\}_{k=1}^K$ (covering nets on S^2), leading to K block LMIs, or (ii) via an S-procedure relaxation: there exists $\lambda \geq 0$ such that

$$\begin{bmatrix} HA^0 & HA(n) \\ A(n)^\top H & HA^0 \end{bmatrix} - \lambda \begin{bmatrix} 0 & 0 \\ 0 & HA^0 \end{bmatrix} (n^\top \gamma\, n - 1) \;\succeq\; 0 \quad \forall n \in \mathbb{R}^3, \tag{11.11}$$

which is conservative but certifies the cone containment with a single LMI in λ and H.

Dissipative terms and GKSL compatibility.

If B contains dissipative/GKSL pieces, impose the H–contractivity LMI

$$\mathcal{D}_H \;:=\; HB + B^\top H \;\preceq\; 0, \tag{11.12}$$

ensuring that dissipation does not spoil well-posed hyperbolic evolution and is compatible with CP/contractivity gates (cf. §4).

Acceptance checklist (cone LMIs).

Checklist 11.2: Cone feasibility (hyperbolicity/causality)

- Solve for $H \succ 0$ satisfying (11.8) on each DoVslab.
- Enforce (11.9) on a covering set $\{n_k\}$ (or S-procedure (11.11)); report the smallest cone margin $\min_k \lambda_{\min}$ of the block matrices.
- (Second order) Verify (11.10) on null/timelike covectors.
- If present, check (11.12) for GKSL pieces in B.
- Flag any violation as a *kill* for Gate G/S; otherwise publish margins and the certificate (H, λ).

Artifact 11.2: AR: Cone-LMI feasibility suite

Implements: (i) construction of frozen-coefficient A^μ on (M, g_{eff}); (ii) SDP for H and block-LMIs (11.9); (iii) principal-symbol PSD tests (11.10); (iv) optional S-procedure relaxation (11.11); (v) regression tests with analytic models saturating the cone.

11.2 Near–Far Detector Consistency; Orthogonal Templates to Break Degeneracies

Scope. We formalize a near–far (ND–FD) mapping that transports spectral information from the near detector to predictions at the far detector using the MSP-consistent oscillation kernel on (M, g_{eff}), and we define *orthogonal template* families that decorrelate flux/cross-section/systematic effects from new-physics deformations (NSI/CPT-odd) on the DoV. Consistency failures are *kill* signatures for the modeling layer (flux → cross section → detector response) or for the assumed physics sector.

11.2.1 ND–FD maps

Event-rate operator form.

Let E denote true neutrino energy and b a reconstructed-energy bin. For a beam flavor α and FD appearance channel β at baseline L_F, the predicted FD spectrum is

$$n_b^F(\vartheta,\nu) \;=\; \int \mathrm{d}E\; \underbrace{R_b^F(E)\,\sigma_\beta(E,\nu)\,\epsilon_\beta^F(E,\nu)}_{\text{FD response}}\; \underbrace{\Phi_\alpha(E,\nu)}_{\text{flux at source}}\; \underbrace{P_{\alpha\beta}(L_F,E;\vartheta)}_{\text{oscillation kernel}} \;+\; b_b^F(\nu), \tag{11.13}$$

where ϑ collects physics parameters (mass splittings/mixings and NP), and ν denotes nuisance parameters (flux, cross sections, detector effects). For the ND (baseline $L_N \ll L_F$), with channel $\alpha \to X$ (typically CC of α),

$$n_a^N(\nu) \;=\; \int \mathrm{d}E\; R_a^N(E)\,\sigma_\alpha(E,\nu)\,\epsilon_\alpha^N(E,\nu)\,\Phi_\alpha(E,\nu) \;+\; b_a^N(\nu), \tag{11.14}$$

with bins a.

Linear transfer map (unfold–refold).

Define the *ND unfolding* operator $\mathcal{U}^N$ that maps n^N to an estimate of the *source flux* $\widehat{\Phi}_\alpha$ (regularized inversion of $R^N\sigma_\alpha\epsilon^N$ on the DoV). The *FD refolding* operator $\mathcal{F}^F(\vartheta)$ maps a source flux to an FD spectrum by applying $P_{\alpha\beta}(L_F,E;\vartheta)$, $\sigma_\beta\epsilon^F$ and R^F:

$$\widehat{n}^F(\vartheta,\nu) \;=\; \mathcal{F}^F(\vartheta,\nu)\,\mathcal{U}^N(\nu)\,n^N(\nu) \;+\; \Delta_{\text{geom}}, \tag{11.15}$$

where Δ_{geom} collects geometry-induced corrections (beamline/solid-angle differences, acceptance mismatches) calibrated on the DoV. In binned form:

$$\widehat{n}_b^F \;=\; \sum_a M_{ba}(\vartheta,\nu)\,n_a^N \;+\; \Delta_b^F, \qquad M_{ba} = \int \mathrm{d}E\; R_b^F\,\sigma_\beta\epsilon^F\,P_{\alpha\beta}(L_F,E;\vartheta)\,\left[\mathcal{U}^N R^N \sigma_\alpha \epsilon^N\right]^{-1}_{Ea}. \tag{11.16}$$

Gate TQ demands that $\mathcal{U}^N$ be built from local, CP-consistent instrument models; Gate S enforces positivity/regularization stability.

Consistency statistic and gates.

Let $\mathbf{r} := n^F - \widehat{n}^F(\vartheta, \nu)$ be the *near–far residual* and Σ_{NF} its covariance including ND/FD statistical errors, systematics, and DoVslacks (colored/cone; cf. §7.3). Define

$$\chi^2_{NF}(\vartheta, \nu) \;:=\; \mathbf{r}^\top \Sigma_{NF}^{-1} \mathbf{r}, \qquad p_{NF} = \Pr\Big(\chi^2 \geq \chi^2_{NF} \,|\, \mathrm{dof}\Big). \tag{11.17}$$

Acceptance. A consistent ND–FD mapping requires p_{NF} above the declared gate threshold on the DoV. Persistent deficits/excesses localized in (L, E) that cannot be absorbed by allowed ν indicate either a physics effect (NP) or a modeling breakdown. We report *pulls* and *eigenmodes* of $\Sigma_{NF}^{-1/2}\mathbf{r}$ to localize the tension.

Joint likelihood with nuisance profiling.

The end-to-end likelihood combining ND and FD is

$$\mathcal{L}(\mathcal{D}_N, \mathcal{D}_F \mid \vartheta, \nu) = \prod_a \mathrm{Pois}\Big(n_a^{N,\mathrm{obs}}; n_a^N(\nu)\Big) \prod_b \mathrm{Pois}\Big(n_b^{F,\mathrm{obs}}; n_b^F(\vartheta, \nu)\Big) \times \mathcal{N}(\nu;\, 0, \Sigma_\nu) \times \mathbf{1}_{\mathrm{gates}}(\vartheta, \nu), \tag{11.18}$$

with Gaussian priors on ν (flux shapes, cross sections, energy scales) and a gate indicator enforcing S, TQ, G. Profiling/marginalization yields posteriors for ϑ; the ND–FD statistic (11.17) is monitored as a consistency metric.

Double ratios and null tests.

Define bin-wise double ratios that cancel flux and cross-section to first order:

$$\mathcal{R}_b \;:=\; \frac{n_b^F / n_{\pi(b)}^N}{\widehat{n}_b^F / \widehat{n}_{\pi(b)}^N}, \tag{11.19}$$

where π pairs ND and FD bins by transport along the best-fit energy map. Under the null (standard oscillations, consistent modeling) one expects $\mathcal{R}_b \simeq 1$ within covariances; coherent departures trace NP templates.

11.2.2 Orthogonal templates

Template dictionary.

We model deformations of the predicted spectra as a linear combination of *templates* $t^{(k)}$ (vectors on the concatenated ND+FD bin space):

$$\delta n \;=\; \sum_{k=1}^{K} \theta_k \, t^{(k)}, \qquad \theta_k \in \mathbb{R} \text{ (small parameters)}. \tag{11.20}$$

Families include (i) *physics*—NSI ($\varepsilon_{e\mu}, \varepsilon_{e\tau}, \varepsilon_{ee}-\varepsilon_{\mu\mu}$), CPT-odd ($a_{e\mu}$), sterile-like oscillatory modulations; (ii) *systematics*—flux tilt/curvature, cross-section shape, energy-scale nonlinearity, efficiency drifts.

Covariance inner product and orthogonalization.

Let Σ be the block covariance of the concatenated ND+FD bins under the null model. Define

$$\langle u, v\rangle_\Sigma \;:=\; u^\top \Sigma^{-1} v, \qquad \|v\|_\Sigma^2 = \langle v, v\rangle_\Sigma. \tag{11.21}$$

Construct an *orthonormal* template set $\{\tilde{t}^{(k)}\}$ via Gram–Schmidt in $\langle\cdot,\cdot\rangle_\Sigma$, ordering physics templates first so that systematics are orthogonalized *against* them. This yields uncorrelated estimators for θ_k at leading order and minimizes leakage of systematics into physics subspace:

$$\widehat{\theta}_k \;=\; \langle \delta n, \tilde{t}^{(k)}\rangle_\Sigma, \qquad \mathrm{Var}(\widehat{\theta}_k) \;=\; 1 \quad (\text{for orthonormal } \tilde{t}^{(k)}). \tag{11.22}$$

Fisher information and degeneracy breaking.

For a differentiable prediction $n(\xi)$ with parameters ξ, the Fisher matrix is $F = J^\top \Sigma^{-1} J$ with $J_{bk} = \partial n_b/\partial \xi_k$. Near degeneracies correspond to small eigenvalues of F. Augmenting the analysis with orthonormal templates adds rows to J (or equivalently pseudo-data with penalties), *lifting* degeneracies and increasing $\det F$. We choose templates to maximize $\Delta \log\det F$ subject to gates (S positivity, DI slacks, DoVsupport).

Matched filters for CP-odd ridges and NSI fringes.

Let t^{CPT} be the signed difference vector predicted by (10.30) between ν and $\bar{\nu}$ spectra across ND/FD (stacked), and let $t^{\mathrm{NSI}}_{e\mu}, t^{\mathrm{NSI}}_{e\tau}$ represent the leading off-diagonal NSI phase-shifted fringes from (10.29). Their Σ-orthonormalized versions provide *matched filters* that are maximally sensitive to these effects and minimally correlated with flux/cross-section tilts.

Nulling templates for standard-oscillation derivatives.

Define derivatives $d_\delta = \partial n/\partial\delta$, $d_{\theta_{13}}$, etc. Construct *nulling* templates $\tilde{t}$ that satisfy $\langle \tilde{t}, d_\delta\rangle_\Sigma = \langle \tilde{t}, d_{\theta_{13}}\rangle_\Sigma = 0$ while aligning with a physics deformation (e.g. t^{CPT}). This prevents bias from imperfect knowledge of standard parameters when extracting NP.

Template priors and DoVsupport.

Systematic templates carry Gaussian priors $\theta_k \sim \mathcal{N}(0, \tau_k^2)$ with τ_k fixed by calibration data; physics templates are unpenalized or carry weak priors derived from MSP geometry (e.g. rank–one NSI). All templates must have support within the declared DoV(energy/baseline ranges; Gate G).

Composite statistic and acceptance.

With orthonormal templates, the composite test statistic decomposes:

$$\chi^2 \;=\; \chi^2_{\perp} \;+\; \sum_{k\in\mathcal{P}} \widehat{\theta}_k^2 \;+\; \sum_{j\in\mathcal{S}} \frac{\widehat{\theta}_j^2}{\tau_j^2}, \tag{11.23}$$

where $\mathcal{P}$ indexes physics templates and $\mathcal{S}$ systematics. *Acceptance* requires: (i) ND–FD χ^2_{NF} within gate; (ii) no single physics template exceeds its discovery threshold unless corroborated across orthogonal channels/baselines; (iii) positivity/causality gates satisfied.

Checklist 11.3: ND–FD & templates acceptance

- Build $\mathcal{U}^N$, $\mathcal{F}^F$ and M (Eq. (11.16)); verify stability/positivity.
- Evaluate χ^2_{NF} (Eq. (11.17)) and eigenmodes of residuals.
- Assemble template dictionary, orthonormalize in $\langle\cdot,\cdot\rangle_\Sigma$ (Eq. (11.21)).
- Quantify degeneracy lift via $\Delta \log\det F$; report variances (11.22).
- Enforce gates $\mathsf{S}, \mathsf{TQ}, \mathsf{G}$ and DI slacks; publish pulls and margins.

Artifact 11.3: AR: ND–FD map & template suite

Implements: (i) regularized unfolding $\mathcal{U}^N$ and FD refolding $\mathcal{F}^F$; (ii) covariance construction Σ_{NF} with colored/cone slacks; (iii) orthonormal template builder (physics & systematics), matched filters, and nulling; (iv) Fisher diagnostics and $\Delta \log\det F$ reporting; (v) regression tests on simulated ND–FD pairs with planted NSI/CPT signals.

11.3 Decision Tables: Which Observable Kills Which Hypothesis; Revision Paths

Scope. We organize falsifiability into *kill windows*—confidence regions in the space of observables that must contain the data on the declared DoV—and a *revision playbook*—a directed, minimal-change procedure to repair a failed gate while preserving MSP (S0–S4). Each window is tied to a specific hypothesis or gate, providing a transparent "observable $\Rightarrow$ hypothesis" map.

11.3.1 Kill windows

General statistical definition.

Let $O \in \mathbb{R}^k$ be a vector of observables with MSP prediction $O_\star(\Theta)$ and total covariance $\Sigma = \Sigma_{\text{geom}} + \Sigma_{\text{run}} + \Sigma_{\text{thr}} + \Sigma_{\text{sys}} + \Sigma_{\text{DI}}$ (geometry, running, threshold/matching, systematics, and device-independent slacks). The *p-level kill window* is the ellipsoid

$$\mathcal{W}_p \;:=\; \Big\{ O \;:\; (O - O_\star)^\top \Sigma^{-1} (O - O_\star) \;\le\; \chi^2_{p,k} \Big\}. \tag{11.24}$$

Kill rule: if the measured $\widehat{O} \notin \mathcal{W}_p$, the associated hypothesis is rejected on the DoV(gate failure).

Single-observable windows.

For $k = 1$ with prediction $O_\star \pm \sigma$, the window is the interval $[O_\star - \kappa_p \sigma,\; O_\star + \kappa_p \sigma]$ (e.g. $\kappa_{0.95} = 1.96$). We will use this for $\sin^2\theta_W(M_Z)$, $\alpha_3(M_Z)$ envelopes, etc.

Floors and one-sided windows.

For EDM floors (§10.2.2) the theory predicts a *lower bound* $|d|_{\text{floor}}$. With an experimental upper limit $|d|^{95\%}_{\text{exp}}$:

$$|d|^{95\%}_{\text{exp}} \;<\; |d|_{\text{floor}}\,(1 - \delta_{\text{had}}) \quad\Longrightarrow\quad \textit{kill}\ \text{(irreducible CP floor violated)}. \tag{11.25}$$

Here δ_{had} accounts for conservative hadronic/nuclear matrix–element uncertainties.

Forward-limit positivity windows.

Let $\lambda^{(2k)}_{\min}$ denote the smallest eigenvalue of the Gram–projected Wilson matrices $\mathcal{C}^{(2k)}$ (§11.1.1). The window is $\lambda^{(2k)}_{\min} \ge 0$ for all tested channel/polarization blocks; any $\lambda^{(2k)}_{\min} < -\epsilon_{\text{num}}$ is a *kill* (Gate S).

Cone-LMI windows.

Let $\mu_{\min}$ be the minimum (over sampled directions n on the unit sphere) of the smallest eigenvalue of the block matrix in (11.9); the window is $\mu_{\min} \ge 0$ with margin. Negative $\mu_{\min}$ *kills* the hyperbolicity/causality gate (Gates G/S).

ND–FD consistency window.

With the residual vector $\mathbf{r}$ and covariance Σ_{NF} (§11.2.1), define $\chi^2_{NF} = \mathbf{r}^\top \Sigma^{-1}_{NF} \mathbf{r}$. The window is $p_{NF} > p_{\text{gate}}$ for the stated degrees of freedom; otherwise *kill* (modeling or physics inconsistency).

Priority and dominance.

If multiple windows fail, we adopt the dominance order:

$$\text{Cone/positivity} \;\Rightarrow\; \text{Forward-limit} \;\Rightarrow\; \text{ND–FD} \;\Rightarrow\; \text{Precision/EDM} \;\Rightarrow\; \text{Gauge geometry/NSI}.$$

Hypothesis / Gate	Observable(s)	Test	Kill window (fail condition)	Notes
Born-limit (B)	$p_i = \mathrm{Tr}(\Pi_i \rho_\infty)$ (long-time stats)	χ^2 or KS	Deviations outside $\mathcal{W}_p$ after finite-time bounds (§3.3)	Requires calibrated POVMs
No-signalling (G)	DI bounds (CHSH, DI regions)	DI feasibility	Outside DI region under colored slacks (§7.3)	Uses beyond-Markov embedding
Gauge geometry (G*)	$\alpha_{1,2,3}(\mu_{\mathrm{ref}})$, $\sin^2\theta_W(\mu_{\mathrm{ref}})$	Ellipsoid	$\widehat{O} \notin \mathcal{W}_p$	From Lie–metric (§9.1)
EW prediction	$\sin^2\theta_W(M_Z)$, $\alpha_3(M_Z)$	Ellipsoid	$\widehat{O} \notin \mathcal{W}_p$	With RG/thresholds (§10.1.2)
Oblique (TQ)	(S, T, U)	Ellipsoid	$\widehat{O} \notin \mathcal{W}_p$	Polarizations from $\delta\mathsf{G}$ (§10.2.1)
EDM floors (CP)	$d_e, d_n, \ldots$	One-sided	$\lvert d\rvert^{\mathrm{UL}} < \lvert d\rvert_{\mathrm{floor}}(1-\delta_{\mathrm{had}})$	Barr–Zee floor (§10.2.2)
NSI/CPT	$P_{\alpha\beta}(L, E)$ shapes	Likelihood	Orthogonal template pulls $\gg$ thresholds	Uses matched filters (§11.2.2)
Forward-limit	$\mathcal{C}^{(2k)}$	LMI	$\lambda_{\min}^{(2k)} < 0$	Dispersion gate (§11.1.1)
Cone causality	Block LMI $\mu_{\min}$	LMI	$\mu_{\min} < 0$	Gravity cone dominance (§11.1.2)
ND–FD map	n^N, n^F	χ^2_{NF}	$p_{NF} < p_{\mathrm{gate}}$	Near–far consistency (§11.2.1)

Table 11.1: Decision windows: observables, tests, and kill conditions on the declared DoV.

Earlier (more kinematic) gates take precedence; later (model-dependent) windows are reassessed only after kinematic repairs.

11.3.2 Revision playbook

Decision graph (minimal repair first).

Given a failed window, traverse the following DAG until the gate is satisfied or MSP is abandoned for the current DoV:

1. **Validate instruments and DoV.** Recheck calibration/Hadamard assumptions and pointer POVMs (H5); shrink DoVif necessary (colored/noise, locality radius L_{loc}).

2. **Kinematic repairs (non-parametric).**

   - If *cone LMI* fails: search symmetrizer H (SDP); if infeasible, adjust principal symbol (reweight kinetic kernel G within uncertainties) or tighten DoV.
   - If *forward-limit* fails: remove offending higher-derivative operator combinations ($\mathcal{C}^{(2k)} \succeq 0$) or relocate thresholds so dispersion prerequisites hold.

3. **Spectral/microscopic revisions.**

   - *Thresholds/matching:* move DoV boundaries M_t, recompute finite matching $(\delta^{\mathrm{fin}}, \Delta^{\mathrm{fin}})$ from CP embeddings.

- *Colored embeddings:* add pseudomodes to restore CP and locality; retune filter functions (Chap. 5).

4. **Geometric kernel updates.** Refit the fiber Lie–metric weights w (S0) within prior ranges; update G and abelian mixing; repropagate to M_Z (§10.1.2).

5. **Overlap sector adjustments.** Modify profile modes $\{\Xi_m\}$ and their coefficients y_m (§9.3.1) under symmetry constraints to fix Yukawa textures, then recompute CKM/PMNS relations and EDM floors.

6. **Template/likelihood refresh.** Rebuild orthogonal templates and ND–FD maps with updated kernels; reprofile nuisances and recompute χ^2_{NF} and physics pulls.

7. **Escalation (structural).** If all above fail with margins, consider (i) changing residual fiber symmetry $H \subset G$, (ii) altering matter representation assignments, or (iii) revising MSP* inputs $\{\kappa_G, \mathcal{N}_u, \mathcal{N}_d, \mathcal{N}_\ell\}$—documented as a new branch with fresh acceptance gates.

Termination and provenance.

Each revision must pass *all* earlier gates (S/G/TQ) and publish *provenance*: diffs of G, thresholds, overlaps, and the resulting deltas in all windows with updated covariances.

Artifact 11.4: Artifacts. Public cards/seeds/HDF5; regression tests

Deliver: (i) **Decision tables** (CSV/TeX) listing windows, inputs, and pass/fail with margins; (ii) **Seeds/HDF5** packs containing $\mathsf{G}(\mu_{\text{ref}})$, abelian kernels, RG trajectories, overlap vectors, threshold sets, and cone/forward certificates (H, λ); (iii) **Public cards** summarizing predictions and windows (§10); (iv) **Regression tests** that re-run the full pipeline and assert gate satisfaction, including randomized perturbations within declared covariances.

Part VI

Tabletop Quantum-Control Probes

Chapter 12

Colored-Zeno Protocols and Noise Spectroscopy

Prelude. On the declared DoV, the spectral medium supports *colored* system–meter noise with finite correlation times (Chap. 5). Control fields can *Zeno-engineer* the dynamics by shaping the frequency-domain sensitivity of the system to that noise. This chapter builds a convex (LMI) design formalism for filter functions that (i) suppress deleterious bands, (ii) enforce causal/positivity gates, and (iii) export reproducible reference sequences and code.

12.1 Filter-function Design and Robustness (Convex LMIs); Reference Sequences/Code

System, noise, and toggling frame.

Let $H(t) = H_S + H_c(t) + H_{SB}$ with $H_{SB} = \sum_\alpha S_\alpha \otimes B_\alpha$, and control $H_c(t) = \sum_k u_k(t)\, C_k$ bounded and piecewise continuous on $[0, T]$. In the toggling frame $U_c(t) = \mathcal{T} \exp\!\Big(-i \int_0^t H_c\Big)$, the couplings become $\tilde{S}_\alpha(t) = U_c^\dagger(t) S_\alpha U_c(t)$. For dephasing about Z, $\tilde{S}(t) = y(t)\, Z$ with a *switching function* $y(t) \in \{\pm 1\}$ (for instantaneous π-pulses) or $|y(t)| \le 1$ (bounded drive).

Filter functions and colored spectra.

Let $S_{\alpha\beta}(\omega)$ be the one-sided bath spectral density matrix on the DoV. Define the (matrix) filter

$$\mathsf{F}_{\alpha\beta}(\omega; T) \;:=\; \tilde{y}_\alpha(\omega)\, \tilde{y}_\beta(\omega)^*, \qquad \tilde{y}_\alpha(\omega) := \int_0^T \mathrm{d}t\; e^{i\omega t}\, y_\alpha(t). \tag{12.1}$$

In the weak-coupling/second-cumulant regime consistent with Gate S, the leading infidelity (or dephasing variance) is

$$\Phi(T) \;\approx\; \frac{1}{2\pi} \int_0^\infty \mathrm{d}\omega\; \mathrm{Tr}\Big[\mathsf{S}(\omega)\, \mathsf{F}(\omega; T)\Big] \;\equiv\; \langle \mathsf{S}, \mathsf{F} \rangle, \tag{12.2}$$

where $\langle\cdot,\cdot\rangle$ denotes the L^2 inner product over frequency. For single-axis dephasing, $\Phi(T) = \frac{1}{2\pi}\int S(\omega)\,|Y(\omega)|^2\,\mathrm{d}\omega$.

Design objective (weighted suppression).

Given a nonnegative weight $W(\omega)$ capturing task priorities (e.g. notches at pseudomode peaks from Chap. 5), we minimize

$$J(y) \;:=\; \frac{1}{2\pi}\int_0^\infty \mathrm{d}\omega\; W(\omega)\,\mathrm{Tr}\big[\mathsf{F}(\omega;T)\big] \quad \text{subject to hardware and gate constraints.} \tag{12.3}$$

When a parametric bath $S(\omega;\theta)$ is available, replace $W(\omega)$ by $\mathbb{E}_\theta[S(\omega;\theta)]$ or adopt a minimax envelope.

Basis expansion and Gram lifting (convexification).

Expand each switching function in a fixed time basis $\{b_m(t)\}_{m=0}^{M-1}$ (e.g. piecewise-constant Walsh, Fourier, or spline atoms):

$$y_\alpha(t) \;=\; \sum_{m=0}^{M-1} a_{\alpha m}\, b_m(t), \qquad \tilde{y}_\alpha(\omega) \;=\; \sum_m a_{\alpha m}\,\tilde{b}_m(\omega). \tag{12.4}$$

Then $\mathsf{F}(\omega;T)$ is quadratic in the real coefficients $a_{\alpha m}$. Introduce the Hermitian Gram matrix

$$G_{\alpha m;\,\beta n} \;:=\; a_{\alpha m}\, a^*_{\beta n} \quad (\text{rank-1},\; G \succeq 0), \tag{12.5}$$

and precompute the positive semidefinite *kernel* matrices

$$Q_{(\alpha m),(\beta n)} \;:=\; \frac{1}{2\pi}\int_0^\infty \mathrm{d}\omega\; W(\omega)\,\tilde{b}_m(\omega)\,\tilde{b}_n(\omega)^*\,\delta_{\alpha\beta}. \tag{12.6}$$

The cost becomes *linear* in G:

$$J \;=\; \sum_\alpha \sum_{m,n} Q_{mn}\, G_{\alpha m;\,\alpha n} \;=\; \mathrm{Tr}(Q\,G), \tag{12.7}$$

and the design problem reads

$$\min_{G\succeq 0}\; \mathrm{Tr}(Q\,G) \quad \text{subject to linear constraints on } G \quad (\text{rank-1 dropped: convex LMI relaxation}). \tag{12.8}$$

A near-rank-1 optimizer yields coefficients via principal component extraction.

Linear constraints (hardware, notches, smoothness).

The following *affine* constraints in a are linear in G after lifting:

1. **Amplitude/energy:** $\int_0^T y_\alpha(t)^2\,\mathrm{d}t \le P_\alpha \Rightarrow \mathrm{Tr}(R\,G) \le P_\alpha$ with $R_{mn} = \int b_m b_n$.
2. **DC/AC nulls:** $Y_\alpha(0) = \sum_m a_{\alpha m}\int b_m = 0$ and $Y_\alpha(\omega_0) = 0$; more generally $\partial^r_\omega \tilde{y}_\alpha(\omega_0) = 0$ for $r = 0,\dots,R_*$ impose linear equations in a; lift as $A\,\mathrm{vec}(G) = 0$.

3. **Duty-cycle and slew:** $\int_0^T |y_\alpha(t)|\,\mathrm{d}t \le D_\alpha$ (via convex surrogates on coefficients), and $\int_0^T \dot{y}_\alpha(t)^2 \le S_\alpha$ ($\Rightarrow$ quadratic form $\mathrm{Tr}(S\,G) \le S_\alpha$).

4. **Instantaneous-pulse timing:** for $y \in \{\pm 1\}$ with N flips at times $0 < t_1 < \cdots < t_N < T$, parameterize in a staircase basis and constrain minimum separation $t_{j+1} - t_j \ge \tau_{\min}$ (linear banded inequalities on the knot vector).

Nonnegativity of trigonometric polynomials (global guarantees).

For piecewise-constant $y(t)$ with periodized extension, $|Y(\omega)|^2$ is a trigonometric polynomial in ω. Requiring band-limited smallness can be phrased as $\Pi(\theta) \ge 0$ for all θ with Π a polynomial in $e^{i\theta}$; Fejér–Riesz/SOS yields an *equivalent Toeplitz LMI* $T(\Pi) \succeq 0$. This furnishes *global* (grid-free) frequency guarantees.

Colored-Zeno constraints.

To enforce Zeno suppression near pseudomode peaks $\{\Omega_\ell\}$ with widths $\{\Gamma_\ell\}$ (Chap. 5), add

$$|Y(\Omega_\ell)|^2 \;\le\; \epsilon_\ell, \qquad \int_{\Omega_\ell-\Gamma_\ell}^{\Omega_\ell+\Gamma_\ell} |Y(\omega)|^2\,\mathrm{d}\omega \;\le\; \eta_\ell, \tag{12.9}$$

both convex in G via precomputed band kernels.

Multi-axis and matrix filters.

For $\mathsf{S}(\omega) \in \mathbb{C}^{p\times p}$ and p modulated axes, stack the coefficients into a block Gram $G \in \mathrm{SS}_+^{pM}$ and use $J = \mathrm{Tr}(\mathcal{Q}\,G)$ with $\mathcal{Q} = \mathrm{diag}(Q, \ldots, Q)$ or a cross-axis kernel that captures S's off-diagonals.

Reference sequences and realizations.

The LMI machinery recovers and *generalizes* classics: (i) CPMG/CP for narrowband dephasing (zeros at DC); (ii) UDD for algebraic suppression near $\omega = 0$ (enforcing R_* derivatives); (iii) Walsh/DD with hardware-friendly bang–bang; (iv) modulated continuous drives ("cos/sin carriers" in the basis (12.4)) that carve movable notches at Ω_ℓ. We publish a small catalog with their (S, F) footprints.

Robustness via LMIs (preview).

Manufacturing and calibration errors (amplitude miscalibration δA, timing jitter δt, detuning $\delta\Delta$) are modeled by affine perturbations $a \mapsto a + \sum_r \delta_r\, a^{(r)}$; the worst-case cost is $\max_{\delta\in\mathcal{B}} \mathrm{Tr}(Q\,G(\delta))$, upper-bounded by an LMI (S-procedure). Details and margins appear in §12.1.2.

Acceptance and gates.

Gate S: $Q \succeq 0$, $G \succeq 0$, and Toeplitz LMIs certify positivity and global suppression; *Gate* G: bandwidth and drive limits respect the emergent cone; *Gate* TQ: the control embeds into CP channels used in Chap. 4. The design exports (i) the optimal coefficients a, (ii) certified margins (minimal eigenvalues), and (iii) code to reproduce kernels and

sequences.

12.1.1 Design toolchain

Inputs and targets.

Given a horizon T, hardware limits (max amplitude/energy, slew, duty), desired notches $\{(\Omega_\ell, \Gamma_\ell)\}$ from colored pseudomodes, and either a weight $W(\omega) \geq 0$ or a parametric bath $S(\omega; \theta)$ on the DoV, the goal is to synthesize switching functions $y_\alpha(t)$ that minimize the weighted filter cost (12.3) subject to gates.

Basis selection and precomputation.

Choose atoms $\{b_m(t)\}_{m=0}^{M-1}$ (Walsh/piecewise-constant, Fourier, splines), form $\tilde{b}_m(\omega) = \int_0^T e^{i\omega t} b_m(t)\, dt$, and assemble:

$$R_{mn} = \int_0^T b_m(t) b_n(t)\, dt, \quad Q_{mn} = \frac{1}{2\pi} \int_0^\infty W(\omega)\, \tilde{b}_m(\omega) \tilde{b}_n^*(\omega)\, d\omega,$$

plus band kernels $B_{mn}^{(\ell)} = \frac{1}{2\pi} \int_{\Omega_\ell - \Gamma_\ell}^{\Omega_\ell + \Gamma_\ell} \tilde{b}_m\, \tilde{b}_n^*\, d\omega$, and moment vectors $c_m^{(r)}(\omega_0) = \partial_\omega^r \tilde{b}_m(\omega_0)$ for derivative nulls.

Convex lifted program (SDP).

Expand $y_\alpha(t) = \sum_m a_{\alpha m} b_m(t)$, lift to the Gram matrix $G_{\alpha m;\, \beta n} = a_{\alpha m} a_{\beta n}^* \succeq 0$ (rank-1 dropped). Solve

$$\min_{G \succeq 0} \ \mathrm{Tr}(Q\, G) \quad \text{s.t.} \quad \begin{cases} \mathrm{Tr}(R\, G) \leq P_\alpha & \text{(energy)} \\ \sum_{m,n} c_m^{(0)}(0) c_n^{(0)*}(0)\, G_{\alpha m;\, \alpha n} = 0 & \text{(DC null)} \\ \sum_{m,n} c_m^{(r)}(\Omega_\ell) c_n^{(r)*}(\Omega_\ell)\, G_{\alpha m;\, \alpha n} = 0 \quad (r = 0..R_*) & \text{(spectral zeros)} \\ \mathrm{Tr}(B^{(\ell)}\, G) \leq \eta_\ell & \text{(band notches)} \\ \text{linear duty/slew constraints on } G & \text{(hardware)} \end{cases} \tag{12.10}$$

For multi-axis modulation stack blocks so $G \in \mathrm{SS}_+^{pM}$ and use block-diagonal Q or cross-axis kernels matching $\mathsf{S}(\omega)$.

Coefficient recovery and realization.

Extract a from the principal component of G (near-rank-1), optionally refine by a short nonconvex projection that enforces discrete $\{\pm 1\}$ pulses under a trust-region. Emit $y_\alpha(t)$, $Y_\alpha(\omega)$, and certified constraints values $\{\mathrm{Tr}(QG), \mathrm{Tr}(B^{(\ell)}G), \ldots\}$.

Verification and export.

Verify *ex-post* that Toeplitz/SOS nonnegativity for trigonometric polynomials holds when global guarantees are requested; produce HDF5 "seeds" (basis, Q, R, $B^{(\ell)}$, recovered a, G) and reference code to reproduce kernels and plots. Gate compliance ($\mathsf{S}, \mathsf{G}, \mathsf{TQ}$) is attached

as a certificate.

12.1.2 Robustness margins

Error models.

We consider small affine perturbations around the designed a: (i) amplitude miscalibration $1+\varepsilon_A$, (ii) timing jitter δt (zero-mean, var σ_t^2), (iii) detuning Δ (carrier mismatch), (iv) notch-center drift $\delta\Omega_\ell$, (v) control-noise PSD $S_A(\omega)$ (multiplicative). They induce *quadratic* corrections to the cost and constraints that can be upper-bounded by additional PSD kernels.

Quadratic robustification (first-order).

For small jitter/detuning,

$$Y(\omega+\Delta) \;\approx\; Y(\omega) - \Delta\,\partial_\omega Y(\omega) + \tfrac{\Delta^2}{2}\partial_\omega^2 Y(\omega), \quad \mathbb{E}_{\delta t}\Big[|Y(\omega)|^2\Big] \;\approx\; |Y|^2 + \sigma_t^2\,|\partial_\omega Y|^2.$$

Hence augment the objective with derivative kernels

$$Q^{(\partial)}_{mn} \;:=\; \frac{1}{2\pi}\int_0^\infty W(\omega)\,\partial_\omega\tilde{b}_m(\omega)\,\partial_\omega\tilde{b}_n^*(\omega)\,d\omega,$$

and use the *robust* kernel

$$Q_{\text{rob}} \;=\; Q \;+\; \sigma_t^2\,Q^{(\partial)} \;+\; \Delta_{\max}^2\,Q^{(\partial)} \;+\; Q^{(A)}, \qquad Q^{(A)}_{mn} = \frac{1}{2\pi}\int_0^\infty S_A(\omega)\,\tilde{b}_m\tilde{b}_n^*\,d\omega, \tag{12.11}$$

so that $J_{\text{rob}} = \mathrm{Tr}(Q_{\text{rob}}G)$ upper-bounds the worst-case cost on the uncertainty set.

Robust band/zero constraints.

Account for notch-center drift by widening bands to $[\Omega_\ell - (\Gamma_\ell + \delta\Omega_\ell),\, \Omega_\ell + (\Gamma_\ell + \delta\Omega_\ell)]$ (i.e. $B^{(\ell)} \mapsto B^{(\ell)}_{\text{wide}}$), and enforce derivative caps at centers:

$$\sum_{m,n} c_m^{(1)}(\Omega_\ell)c_n^{(1)*}(\Omega_\ell)\,G_{\alpha m;\,\alpha n} \;\le\; \kappa_\ell,$$

so small $\delta\Omega_\ell$ cannot spoil deep notches.

Margins (reported certificates).

We publish the following nonnegative margins:

$$\begin{aligned}
&m_{\text{pow}} = P_\alpha - \mathrm{Tr}(R\,G), \qquad m^{(\ell)}_{\text{band}} = \eta_\ell - \mathrm{Tr}(B^{(\ell)}_{\text{wide}}\,G),\\
&m_{\text{zero}} = \epsilon_0 - \sum_{m,n} c_m^{(0)}(0)c_n^{(0)*}(0)\,G_{\alpha m;\,\alpha n}, \qquad m^{(\ell)}_{\partial} = \kappa_\ell - \sum_{m,n} c_m^{(1)}(\Omega_\ell)c_n^{(1)*}(\Omega_\ell)\,G_{\alpha m;\,\alpha n},\\
&m_J = J_{\text{tol}} - \mathrm{Tr}(Q_{\text{rob}}G).
\end{aligned}$$

Acceptance on the DoVrequires all margins ≥ 0 with declared buffers.

S-procedure (ellipsoidal uncertainties).

For an ellipsoid $\{\delta : \ \delta^\top \Pi\, \delta \le 1\}$ acting linearly on coefficients $a \mapsto a + \sum_r \delta_r a^{(r)}$, the robust inequality $\mathrm{Tr}(Q\, G(\delta)) \le J_{\mathrm{tol}}$ is implied by the existence of $\lambda \ge 0$ such that the quadratic form in δ is dominated:

$$\begin{bmatrix} J_{\mathrm{tol}} - \mathrm{Tr}(QG) - \lambda & \frac{1}{2} q^\top \\ \frac{1}{2} q & \Pi - Q_2 \end{bmatrix} \succeq 0, \qquad q_r = \mathrm{Tr}(Q\, G_r), \quad (Q_2)_{rs} = \mathrm{Tr}(Q\, G_{rs}), \tag{12.12}$$

where G_r, G_{rs} are the first/second differentials of G w.r.t. δ_r; this furnishes a conservative LMI certificate for worst-case cost.

Simulation cross-checks.

We validate robust margins by (i) GKSL Monte Carlo with colored noise consistent with Chap. 5, (ii) randomized hardware perturbations within the uncertainty set, and (iii) empirical extraction of effective filter functions from measured spectra; the empirical cost should lie below J_{rob} with the reported buffers.

Artifact 12.1: AR: Toolchain & robustness suite

Packs: kernels (Q, $Q^{(\partial)}$, $B^{(\ell)}$, R, $c^{(r)}$), SDP configs, recovered (a, G), robustified Q_{rob}, and margin reports $\{m_{\mathrm{pow}}, m^{(\ell)}_{\mathrm{band}}, m_{\mathrm{zero}}, m^{(\ell)}_{\partial}, m_J\}$, plus scripts to reproduce frequency-domain plots and GKSL simulations.

12.2 Parameter-to-Knob Maps: Cold Atoms, Waveguides, (Optional) Superconducting Circuits

Prelude. On the declared DoV, control protocols from §12.1 are realized by experimental *knobs* that shape the system–meter Hamiltonian and the noise spectral density seen by the pointer sectors. We formalize a map

$$\mathcal{M} : \ (\Theta_{\mathrm{th}},\, a) \ \longmapsto \ u \ \longmapsto \ \big(H(t; u),\ \mathsf{S}(\omega; u)\big),$$

where Θ_{th} are *theory* targets (pseudomode centers/widths $\{\Omega_\ell, \Gamma_\ell\}$, weights $\{\lambda_\ell\}$, locality radius, cone gates), a are filter-design coefficients (e.g. from the SDP in (12.8)), and u are hardware knobs. We give explicit, dimensionally consistent formulas for cold atoms and waveguides, plus linear sensitivities $\partial \mathsf{S}/\partial u$ to support robust margins.

12.2.1 Cold-atom implementations

Knob vector.

We collect laser/trap/magnetic knobs in

$$u_{\text{cold}} = \Big(V_0,\ a_{\text{lat}},\ \Omega_R,\ \Delta_R,\ \delta,\ B,\ \omega_{\text{trap}},\ I,\ \phi(t)\Big),$$

where $V_0 = sE_R$ is lattice depth ($E_R = \hbar^2 k^2/2m$), $a_{\text{lat}} = \pi/k$ lattice spacing, Ω_R, Δ_R single-photon Rabi and detuning, δ two-photon detuning, B magnetic field (Feshbach), ω_{trap} harmonic confinement, I laser intensity, and $\phi(t)$ programmable phase (for phase modulation).

System Hamiltonian (effective spin–boson).

For a pseudo-spin encoded in two hyperfine states with Raman coupling,

$$H_S(t) = \frac{\hbar}{2}\Big[\Delta(t)\,\sigma_z + \Omega(t)\,\sigma_x\Big], \quad \Delta(t) = \delta + \delta_{\text{ac}}(I, \Delta_R), \quad \Omega(t) \approx \frac{\Omega_1\Omega_2}{2\Delta_R}, \tag{12.13}$$

with δ_{ac} the differential AC Stark shift. Coupling to a bosonic bath (phonons of a BEC, cavity-assisted modes, or technical phase noise)

$$H_{SB} = \sigma_z \sum_k g_k(u_{\text{cold}})\,(b_k + b_k^\dagger), \qquad H_B = \sum_k \hbar\omega_k(u_{\text{cold}})\, b_k^\dagger b_k, \tag{12.14}$$

yields the spectral density $J(\omega) = \sum_k |g_k|^2 \delta(\omega - \omega_k)$, and the (one-sided) noise PSD $S(\omega)$ entering (12.2) is $S(\omega) \propto J(\omega)\coth(\hbar\omega/2k_BT)$ (on the DoV).

Engineering colored Lorentzians (pseudomodes).

Two practical routes:

(1) *Cavity-assisted bath.* With a single lossy mode (frequency ω_c, linewidth κ_c) dispersively coupled to the atom with rate g_c,

$$S(\omega; u_{\text{cold}}) \ \simeq\ \lambda\,\frac{\kappa_c}{(\omega-\omega_c)^2 + (\kappa_c/2)^2}, \qquad \lambda \propto g_c^2, \quad \kappa_c = \frac{\pi c}{\mathcal{F}L}, \tag{12.15}$$

with cavity length L and finesse $\mathcal{F}$. Knobs: $\omega_c = \pi q c/L$ (mirror piezos), κ_c via mirror coatings (fixed) or intra-cavity loss control, g_c via atomic position (standing-wave antinode).
(2) *Phononic bath in a BEC.* Bogoliubov dispersion $\omega_k = \sqrt{c_s^2 k^2 + (\hbar k^2/2m)^2}$ with sound speed $c_s = \sqrt{gn_0/m}$, coupling $\propto a_s$,

$$g = \frac{4\pi\hbar^2 a_s}{m}, \qquad a_s(B) = a_{\text{bg}}\Big(1 - \frac{\Delta_B}{B - B_0}\Big),$$

so the low-ω slope of $S(\omega)$ is tuned by $a_s(B)$ and density n_0. Close to engineered

rotons/band edges (via lattice V_0) one obtains shaped, peaked spectra.

Filter-function realization (pulse & continuous).

A bang–bang sequence with π-pulses at $\{t_j\}$ realizes $y(t) = (-1)^{N(t)}$ and *approximately* enforces the zeros/constraints of §12.1. Finite pulse width τ_π and Rabi cap $\Omega_{\max}$ enter via a convolution kernel h_τ:

$$Y_{\text{lab}}(\omega) \;\approx\; H_\tau(\omega)\, Y_{\text{ideal}}(\omega), \qquad |H_\tau(\omega)| \;\approx\; \left|\frac{\sin(\omega\tau_\pi/2)}{\omega\tau_\pi/2}\right|. \tag{12.16}$$

For continuous drives, set $\Omega(t) = \Omega_0 + \sum_q \Omega_q \cos(\omega_q t + \varphi_q)$, so $Y(\omega)$ inherits notches at $\{\omega_q\}$; Ω_q and φ_q are direct knobs via AOM/EOM amplitude and phase.

Sensitivity Jacobians (for robustness).

Local linearization supplies

$$\frac{\partial\omega_c}{\partial L} = -\frac{\omega_c}{L}, \quad \frac{\partial\kappa_c}{\partial\mathcal{F}} = -\frac{\kappa_c}{\mathcal{F}}, \quad \frac{\partial c_s}{\partial a_s} = \frac{1}{2}\frac{c_s}{a_s}, \quad \frac{\partial S(\omega)}{\partial a_s} \propto \frac{S(\omega)}{a_s},$$

and $\partial Y/\partial\tau_\pi$ from (12.16). These feed the robust kernels in (12.11).

Gate checks & calibration.

Gate S*:* verify CP contractivity by GKSL fits of idle evolution (Chap. 4); *Gate* G*:* ensure band-limited control obeys the cone (maximal group velocity in the optical lattice); *Gate* TQ*:* instrument POVMs calibrated (H5). Publish HDF5 seeds: $\{u_{\text{cold}}, S(\omega), Y(\omega), \text{margins}\}$.

12.2.2 Waveguide implementations

Knob vector.

For a quantum emitter (QE) coupled to a 1D/2D photonic waveguide or crystal,

$$u_{\text{wg}} = \Big(\omega_0,\ g,\ Z(\omega),\ \omega_f,\ \kappa_f,\ v_g(\omega),\ \beta(\omega),\ \Phi(t)\Big),$$

where ω_0 QE transition, g light–matter coupling, $Z(\omega)$ engineered impedance (Purcell filter), (ω_f, κ_f) filter resonance and linewidth, v_g group velocity, $\beta(\omega)$ dispersion, and $\Phi(t)$ Stark/phase modulation.

Bath shaping and spectral density.

For a single guided band with DOS $\rho(\omega) = L/(\pi v_g(\omega))$, the waveguide-induced spectral density near the QE is

$$J(\omega; u_{\text{wg}}) \;=\; 2\pi\, |g(\omega)|^2\, \rho(\omega)\, \mathcal{T}(\omega), \qquad \mathcal{T}(\omega) \in [0,1] \tag{12.17}$$

with transmission $\mathcal{T}(\omega)$ set by $Z(\omega)$ or a side-coupled resonator. A Purcell filter with (ω_f, κ_f) yields a Lorentzian $\mathcal{T}(\omega) = \kappa_f^2/[(\omega-\omega_f)^2+(\kappa_f/2)^2]$, hence a colored $S(\omega) \propto J(\omega)$ that implements the pseudomode target (12.15). Band-edge engineering (photonic crystal) modifies $\rho(\omega) \propto \sqrt{\omega - \omega_{\text{edge}}}$ and produces non-Lorentzian tails—useful for testing non-Markov embeddings (Chap. 5).

System Hamiltonian and control embedding.

The QE (two-level or Λ) under a Stark tone $V_{\text{ac}}(t)$ obeys

$$H_S(t) = \frac{\hbar}{2}\Big[\Delta_0 + \delta_{\text{Stark}}(t)\Big]\sigma_z + \frac{\hbar}{2}\Omega_d(t)\,\sigma_x, \qquad \delta_{\text{Stark}}(t) \propto \frac{|E_{\text{ac}}(t)|^2}{\Delta_{\text{ac}}}, \tag{12.18}$$

with Ω_d the drive Rabi rate. Phase/amplitude modulation implements the filter $y(t)$ (bang–bang via π pulses or continuous via tones), and the lab window $H_\tau(\omega)$ follows from the drive bandwidth as in (12.16).

Knobs → PSD map and sensitivities.

For the resonant-filter case,

$$\frac{\partial S(\omega)}{\partial \omega_f} \simeq S(\omega)\,\frac{2(\omega-\omega_f)}{(\omega-\omega_f)^2+(\kappa_f/2)^2}, \qquad \frac{\partial S(\omega)}{\partial \kappa_f} \simeq S(\omega)\,\frac{1}{\kappa_f}\left(1 - \frac{\kappa_f^2/4}{(\omega-\omega_f)^2+(\kappa_f/2)^2}\right),$$

and $\partial S/\partial v_g = -S/v_g$ through $\rho(\omega)$. These gradients parameterize robustified kernels Q_{rob} (§12.1.2).

Zeno notches and DI slacks.

Moving notches are created by setting $\Phi(t)$ (phase) or $\Omega_d(t)$ (amplitude) to track Ω_ℓ, with constraints from G (finite propagation speed v_g) and S (GKSL contractivity). Device-independent slacks from §7.3 are included as additive uncertainties on event rates or as bounds on effective coupling norms extracted by tomography.

Calibration and acceptance.

(i) Extract $S(\omega)$ from noise spectroscopy (Ramsey/Carr–Purcell decay vs. filter), fit to the model in (12.17). (ii) Validate $H_S(t)$ parameters via Rabi/CW scans. (iii) Check positivity/hyperbolicity gates (cone LMIs) for the effective PDE describing field propagation in the line. Publish HDF5 seeds $\{u_{\text{wg}}, S(\omega), Y(\omega)\}$ and margins.

Artifact 12.2: AR: Parameter-to-knob maps

For both platforms we export: (1) symbolic maps $u \mapsto (H, \mathsf{S})$ with units; (2) Jacobians $\partial\mathsf{S}/\partial u$ and $\partial Y/\partial u$ for robustness; (3) code to synthesize $y(t)$ from SDP coefficients and hardware constraints; (4) spectroscopy notebooks that invert $n(t)$ or decoherence curves to $\widehat{S}(\omega)$ for gate compliance.

12.3 Sensitivity Forecasts and Systematics Ledgers; Stopping Rules

Prelude. We provide a self-contained forecasting pipeline for tabletop probes that maps *design choices* (filters, sequences, hardware knobs; §12.1, §12.2) to *sensitivity envelopes* on targeted parameters (e.g. pseudomode $\{\Omega_\ell, \Gamma_\ell, \lambda_\ell\}$, robust costs J, GKSL rates), while accounting for colored-noise structure and gate slacks. A parallel *systematics ledger* defines nuisance priors and correlations. Stopping rules close the loop by specifying when data taking or sequence iterations should terminate on the DoV.

12.3.1 Forecast methodology

Measurement model and Asimov data.

Let $d = \{d_m\}_{m=1}^M$ be a stack of observables (e.g. Ramsey contrasts, CPMG/UDD decays at times $\{T_m\}$, spectral points of $\widehat{S}(\omega)$, tomography summaries). For settings u and physics parameters ϑ (e.g. entries of $\mathsf{S}(\omega)$, GKSL rates, filter coefficients), the prediction is $f_m(\vartheta, u)$ with covariance Σ_d (shot noise, readout infidelity, drift). The *Asimov* dataset is $d_m^{\mathrm{A}} = f_m(\vartheta_0, u_0)$ at a nominal point (ϑ_0, u_0) on the DoV.

Binned Poisson/Gaussian likelihood.

We adopt a Gaussian approximation for continuous contrasts and Poisson for counts; the joint log-likelihood is

$$\log \mathcal{L}(d \mid \vartheta, \nu) = -\tfrac{1}{2}\Big(d - f(\vartheta,\nu)\Big)^\top \Sigma_d^{-1}\Big(d - f(\vartheta,\nu)\Big) \;-\; \tfrac{1}{2}(\nu-\nu_0)^\top \Sigma_\nu^{-1}(\nu-\nu_0) \;+\; \text{const}, \tag{12.19}$$

where ν collects nuisance parameters from the ledger (below) with Gaussian priors $\mathcal{N}(\nu_0, \Sigma_\nu)$. For count channels we use Poisson terms and linearize for Fisher forecasts.

Fisher information and CRLB.

Stack $\xi = (\vartheta, \nu)$. The Fisher matrix at (ϑ_0, ν_0) is

$$F \;=\; J^\top \Sigma_d^{-1} J \;+\; \Pi, \qquad J_{mk} = \frac{\partial f_m}{\partial \xi_k}\bigg|_{\xi_0}, \quad \Pi = \begin{bmatrix} 0 & 0 \\ 0 & \Sigma_\nu^{-1} \end{bmatrix}. \tag{12.20}$$

The profiled covariance for physics parameters is $\mathrm{Cov}(\widehat{\vartheta}) \;\simeq\; \Big(F_{\vartheta\vartheta} - F_{\vartheta\nu}F_{\nu\nu}^{-1}F_{\nu\vartheta}\Big)^{-1}$. Quote $\sigma(\theta_i) = \sqrt{(F_{\mathrm{prof}}^{-1})_{ii}}$ and correlations.

Exposure and scaling laws.

Let the exposure scale as $\mathcal{E}$ (e.g. number of shots $\times$ duty). For shot-limited channels, $\Sigma_d \propto 1/\mathcal{E}$ and $F \propto \mathcal{E}$, hence $\sigma(\theta) \propto \mathcal{E}^{-1/2}$. In presence of *floors* (from calibration, drift, DI

slacks) we model $\Sigma_d = \Sigma_{\text{stat}}/\mathcal{E} + \Sigma_{\text{floor}}$ so sensitivities saturate to a plateau; this informs stopping (below).

Discovery and exclusion surrogates.

For a scalar effect θ with null $\theta = 0$, the Asimov discovery significance is

$$Z_{\text{A}} \approx \frac{|\theta_{\text{alt}}|}{\sigma(\theta)} \quad \text{(profiled Fisher)}, \tag{12.21}$$

and the 95% CL exclusion (one-sided) reach is $|\theta|_{\text{excl}} \approx 1.64\,\sigma(\theta)$. For vectors, use quadratic forms with F_{prof} (noncentral χ^2).

Design-of-experiments (DOE).

For a set of candidate sequences $\{u_r\}_{r=1}^{R}$ (filters, durations, notches), build block-diagonal Σ_d and Jacobians $J^{(r)}$. Choose the portfolio weights $\{w_r\}$ by maximizing a convex information criterion:

$$\max_{w_r \geq 0,\ \sum w_r = 1} \Phi(F(w)), \qquad F(w) = \sum_r w_r\, (J^{(r)})^\top \Sigma_{d,r}^{-1} J^{(r)} + \Pi, \tag{12.22}$$

with Φ one of D-optimal ($\log\det$), A-optimal ($-\operatorname{Tr} F^{-1}$), or E-optimal ($\lambda_{\min}(F)$), subject to gate constraints (bandwidth, power, cone). This yields the time budget per sequence.

Template augmentation.

If the observable map f includes orthonormal *templates* $t^{(k)}$ (§11.2.2) in the measurement stack, append their rows to J so the Fisher matrix directly captures degeneracy lifts between physics and systematic subspaces.

Coverage and LEE.

Validate coverage by toy Monte Carlo at a few support points; adjust Z for a look-elsewhere factor N_{eff} (scan over notch positions or pseudomode frequencies) via $p \mapsto 1 - \big(1-p\big)^{N_{\text{eff}}}$.

12.3.2 Systematics ledger & stopping

Ledger (nuisance catalogue and priors).

We maintain a living *ledger* of nuisances with symbols, priors, and correlations; examples (platform-agnostic core + platform-specific):

- **Readout & SPAM:** state-prep error ϵ_{SP}, measurement infidelity ϵ_{M}; priors from calibration runs.
- **Amplitude/timing:** drive gain drift ε_A, timing jitter std σ_t, finite-width window τ_π (cf. (12.16)).
- **Detuning/drift:** slow $\Delta(t)$ modeled as OU process with diffusion D_Δ; included

via correlated covariance in Σ_d.

- **Bath parameters:** $(\Omega_\ell, \Gamma_\ell, \lambda_\ell)$ priors from spectroscopy; correlations across ℓ if shared hardware origin.
- **PSD leakage:** residual outside designed bands, modeled via Toeplitz LMI slack (nonnegativity polynomial) with tolerance ϵ_{PSD}.
- **DI & cone slacks:** additive rate slack ζ_{col} (colored) and ϵ_{cone} (cone), as in §7.3 and §11.1.2.

Assemble $\nu = (\epsilon_{\mathrm{SP}}, \epsilon_{\mathrm{M}}, \varepsilon_A, \sigma_t, \tau_\pi, \Delta, \{\Omega_\ell, \Gamma_\ell, \lambda_\ell\}, \zeta_{\mathrm{col}}, \epsilon_{\mathrm{cone}}, \ldots)$ with Σ_ν block-structured; propagate into F via (12.20).

Precision, discovery, and kill targets.

Define thresholds:

$$\text{Precision target: } \sigma(\theta_i) \le \tau_i; \quad \text{Discovery: } Z_{\mathrm{A}} \ge Z_*; \quad \text{Kill: } \widehat{O} \notin \mathcal{W}_p \text{ (window (11.24))}.$$

The forecast outputs time/exposure $\mathcal{E}$ required to reach each target under the DOE weights w_r.

Sequential analysis and alpha-spending.

For interim looks at exposures $\{\mathcal{E}_k\}$, use a group-sequential design with O'Brien–Fleming spending:

$$Z_{\mathrm{crit}}^{(k)} \approx \Phi^{-1}\Big(1 - \tfrac{\alpha}{2}\sqrt{\tfrac{\mathcal{E}_K}{\mathcal{E}_k}}\Big),$$

ensuring the overall Type-I error α under repeated peeks (Gate reporting). Alternatively, a Bayesian stop uses posterior odds $\mathrm{BF}_{10} \ge B_*$.

SPRT for binary hypotheses (tabletop kill).

For H_0 (MSP value) vs H_1 (shifted by δ), the cumulative log-likelihood ratio at look k is $\Lambda_k = \sum_m \log \frac{p(d_m|H_1)}{p(d_m|H_0)}$. Stop and *kill* if $\Lambda_k \ge \ln\frac{1-\beta}{\alpha}$; accept H_0 if $\Lambda_k \le \ln\frac{\beta}{1-\alpha}$; otherwise continue. Choose (α, β) per gate policy.

Utility-optimized stopping (cost-aware).

Let $Z(\mathcal{E})$ be the profiled significance or $1/\sigma(\theta)$, and $C(\mathcal{E})$ the resource cost. Choose $\mathcal{E}^\star = \arg\max_{\mathcal{E}} U(\mathcal{E})$, $U = Z(\mathcal{E}) - \lambda C(\mathcal{E})$, or stop when marginal gain falls below threshold: $\frac{\mathrm{d}Z}{\mathrm{d}\mathcal{E}} \le \epsilon_Z$. Include floors by modeling $Z(\mathcal{E})$ saturation.

Ledger-to-stopping hooks.

Certain ledger entries impose *hard stops*: if cone LMI margin $m_{\mathrm{cone}} < 0$ (§11.1.2) or forward-limit LMI fails (§11.1.1), data taking must halt for revision (Gate S/G). If robustness margins (§12.1.2) become negative, re-run the SDP.

Reporting.

For each target θ and platform, publish: (i) $D/A/E$-optimal DOE weights, (ii) $\sigma(\theta)$ vs exposure with floors, (iii) discovery/exclusion curves, (iv) sequential boundaries, (v) the full ledger (ν, Σ_ν) and pull distributions from toys, (vi) gate compliance logs.

Artifact 12.3: AR: Forecast & ledger pack

HDF5 seeds: Asimov d^{A}, Jacobians J, covariances (Σ_d, Σ_ν), DOE weights $\{w_r\}$, sensitivity curves $\sigma(\theta)$ and $Z(\mathcal{E})$, sequential boundaries, and the ledger with correlations. Scripts reproduce Fisher maps, toy coverage, and stopping-time optimization.

Chapter 13

Implementations with Cold Atoms and Waveguides

Prelude. This chapter translates the filter-design formalism of §12.1 and the parameter–to–knob maps of §12.2 into *operational* implementations on two tabletop platforms: (i) cold atoms (hyperfine qubits in optical lattices/cavities/BEC environments), and (ii) quantum emitters embedded in photonic waveguides/crystals. We specify executable *sequences*, their *readouts*, and a minimal set of *calibration/health checks* that certify Gate compliance (S/G/TQ) on the declared DoV. All formulas are dimensionally consistent and reference the filter/kernel objects $\{Y(\omega), \mathsf{F}(\omega), Q, R, B^{(\ell)}\}$ defined earlier.

13.1 Experimental Sequences and Readouts; Calibration/Health Checks

13.1.1 Sequences & readouts

Sequence specification (from filters).

Given a target switching function $y(t) \in [-1, 1]$ on $[0, T]$ (from the SDP in (12.8)), we compile to hardware as:

$$\text{bang–bang:} \quad y(t) = (-1)^{N(t)} \iff \pi\text{-pulses at times } \{t_j\}_{j=1}^{N}, \qquad \text{continuous:} \quad y(t) = \sum_q \alpha_q \cos(\omega_q t + \tag{13.1}$$

with π-pulses realized by resonant $\Omega(t)$ satisfying $\int_{t_j}^{t_j+\tau_\pi} \Omega(t)\, dt = \pi$, and continuous tones programmed via AOM/EOM (cold atoms) or IQ modulator (waveguides). The *lab* filter is $Y_{\text{lab}}(\omega) = H_\tau(\omega)\, Y(\omega)$, with H_τ the finite-width window from (12.16).

Canonical families (cold atoms).

- **Ramsey/Free induction:** $\frac{\pi}{2}$—free—$\frac{\pi}{2}$ with phase ϕ. Contrast $C(T,\phi) = \Re\Big[e^{-\chi(T)}e^{i(\Delta T+\phi)}\Big]$ where

$$\chi(T) \;=\; \frac{1}{2\pi}\int_0^\infty S(\omega)\,|Y_{\rm lab}^{(\rm Ram)}(\omega;T)|^2\,d\omega, \qquad Y^{(\rm Ram)}(\omega) = 2\,\frac{\sin(\omega T/2)}{\omega}\,. \tag{13.2}$$

- **Hahn echo/CPMG(N):** $\frac{\pi}{2}$—$(\tau-\pi-\tau)^N$—$\frac{\pi}{2}$, with $T = 2N\tau$. The ideal $|Y(\omega)|^2$ presents a comb peaked at $\omega_k \simeq (2k-1)\pi/\tau$; band notches/zeros from the SDP are enforced by placing $\{\pi\}$ at the computed $\{t_j\}$.
- **UDD/Walsh:** pulse times $t_j/T = \sin^2\!\Big(\frac{\pi j}{2N+2}\Big)$ (UDD) or Walsh phase flips. These realize derivative zeros at $\omega = 0$ up to order R_* when requested by the constraints.

Canonical families (waveguides).

- **Phase-modulated Stark drive:** $\delta_{\rm Stark}(t) = \sum_q s_q \cos(\omega_q t + \varphi_q)$, implementing $y(t) = \mathrm{sgn}[\cos(\cdot)]$ (bang–bang) or smooth $y(t)$ (tones).
- **Carr–Purcell via π-bursts:** short resonant bursts synthesized by IQ modulation; pulse areas calibrated by Rabi scans (below).

Readouts (cold atoms).

1. **State-dependent fluorescence (SDF):** counts $N_{F=1}, N_{F=2}$ yield $P_{F=2} = N_{F=2}/(N_{F=1}+N_{F=2})$, corrected for SPAM infidelity ($\epsilon_{\rm SP}, \epsilon_{\rm M}$ in the ledger).
2. **Cavity homodyne (if present):** measure quadrature $X_{\rm out}(t)$; its PSD estimates the dephasing rate under a given $y(t)$ via input–output $a_{\rm out} = a_{\rm in} + \sqrt{\kappa}\,\sigma_-$.

Readouts (waveguides).

1. **Heterodyne I/Q:** measure (I, Q) of the transmitted/reflected field; fit to a GKSL lineshape to extract effective rates under the applied filter.
2. **Photon counting:** time-tagging statistics ($\mathrm{g}^{(2)}(\tau)$) under $y(t)$ provide an independent proxy for $\mathsf{S}(\omega)$ features near set notches.

Tomography under colored noise.

We adopt short-depth gate sets (GST-lite) interleaved with filters to reconstruct a time-local GKSL generator $\mathcal{L}_{\rm eff}(t)$ on the DoV, checking CP (contractivity) and locality. The estimator minimizes

$$\min_{\mathcal{L}} \sum_r \|\rho_r^{(\rm meas)} - \mathcal{E}_r[\rho_r^{(\rm in)}]\|_{\Sigma_r}^2 \quad \text{s.t.} \quad \mathcal{E}_r = \mathcal{T}\exp\!\Big(\int \mathcal{L}_{\rm eff}(t)\,dt\Big), \;\; \mathcal{E}_r \text{ CPTP}. \tag{13.3}$$

Violation of CPTP within error bars triggers a DoV shrink or embedding refinement (Chap. 5).

Sequence packs (export).

Each sequence is published with: timing list $\{t_j, \tau_\pi\}$ or tone set $\{\omega_q, \alpha_q, \varphi_q\}$, measured $C(T)$ or spectra, certified notch metrics $\{\mathrm{Tr}(B^{(\ell)}G)\}$, and gate margins.

13.1.2 Calibration & checks

Cold atoms — core routines.

- **Rabi calibration:** scan Ω vs pulse area; enforce $\int \Omega(t)\, dt = \pi$ within tolerance δ_π. Record margin $m_\pi = \pi - \int \Omega$.
- **Detuning & drift:** Ramsey fringe vs T to extract Δ and its OU drift D_Δ; propagate to the ledger and Q_{rob} (§12.1.2).
- **AC Stark shifts:** measure $\delta_{\mathrm{ac}}(I, \Delta_R)$ and update (12.13) so that compiled $y(t)$ targets the intended axis.
- **PSD spectroscopy:** sweep T and sequence family to invert $S(\omega)$ from (13.2); compare to the modeled spectrum from (12.15) or BEC phonons; publish residuals.

Waveguides — core routines.

- **Filter transfer:** measure $\mathcal{T}(\omega)$ via VNA or coherent scan; fit (ω_f, κ_f); compute $J(\omega)$ per (12.17).
- **Group velocity:** extract $v_g(\omega)$ near operating band; ensure cone bandwidth (no superluminal propagation; Gate G).
- **Lineshape GKSL fit:** from heterodyne spectra under $y(t)$, infer effective rates; check CP contractivity (Gate S) and locality window.

Health checks (both platforms).

1. **SPAM audit:** estimate $(\epsilon_{\mathrm{SP}}, \epsilon_{\mathrm{M}})$; update ledger priors and propagate to Fisher forecasts (§12.3).
2. **Cone-LMI certificate:** build frozen-coefficient matrices and verify block LMIs (11.9); publish minimal eigenvalue margin $\mu_{\min}$.
3. **Forward-limit positivity:** assemble low-energy EFT for the probed channels and check $\mathcal{C}^{(2k)} \succeq 0$ (Gate S).
4. **Robustness margins:** recompute $m_{\mathrm{pow}}, m^{(\ell)}_{\mathrm{band}}, m^{(\ell)}_{\partial}, m_J$ (§12.1.2); halt and re-solve SDP if any margin < 0.

Checklist 13.1: Sequence & readout acceptance

- Timing/area calibration within tolerances; detuning drift modeled.
- Measured $S(\omega)$ (spectroscopy) matches modeled within envelopes.
- CPTP tomography within confidence; DI/cone slacks respected.
- Published sequence packs include kernels, margins, and reproducible code.

Artifact 13.1: AR: Implementation packs

For each platform/sequence: `HDF5` seeds (timings or tones, $Y(\omega)$, spectroscopy $\widehat{S}(\omega)$, GKSL fit, margins), scripts to reproduce figures and checks, and CI tests asserting all gates pass on the declared DoV.

13.2 Verification and Tomography (GST/RB under Colored Noise); Error Budgets

Scope. We give operational procedures to (i) *verify* that realized dynamics on the declared DoVadmit a CP/time-local embedding consistent with Gates S and TQ, and (ii) *tomographically* identify effective generators and noise spectra under *colored* (non-Markovian) conditions. We use two complementary tools: Gate Set Tomography (GST) with colored-noise regularization, and Randomized Benchmarking (RB) with filter-function analysis. We close with an explicit, additive *error budget* tying uncertainties to physical and instrumental sources (SPAM, timing, detuning, PSD modeling, DI/cone slacks).

13.2.1 GST/RB under color

Colored-GST model.

Let $\mathcal{G} = \{\,\mathcal{G}_i\,\}_{i=1}^{K}$ be the target gate set on a two-level (or few-level) pointer sector. Under colored noise, the physical map realized by a length-L gate string $s = (i_1, \ldots, i_L)$ is

$$\mathcal{E}_s \;=\; \mathcal{T}\exp\Bigl(\int_0^{T_s} \mathcal{L}_{\text{eff}}(t)\,\mathrm{d}t\Bigr) \;\approx\; \prod_{\ell=1}^{L} \exp\bigl(\Delta t_\ell\,\mathcal{L}_{i_\ell}(t_\ell)\bigr), \qquad \mathcal{L}_i(t) = \mathcal{L}_i^{(0)} + \delta\mathcal{L}_i(t), \tag{13.4}$$

where $\delta\mathcal{L}_i(t)$ is a *colored* fluctuation with spectrum $\mathsf{S}(\omega)$ determined by the platform (§12.2). A convenient CP *embedding* uses a finite pseudomode state $\xi \in \mathbb{R}^q$:

$$\xi_{\ell+1} \;=\; A\,\xi_\ell + \eta_\ell, \qquad \mathcal{L}_{i_\ell}(t_\ell) \;=\; \mathcal{L}_{i_\ell}^{(0)} + \sum_{r=1}^{q} \xi_\ell^{(r)}\,\mathcal{K}_r, \tag{13.5}$$

with $|\rho(A)| < 1$ (mixing gap; Gate S) and η_ℓ zero-mean with covariance fixed by S. The observable probability for a SPAM trio (ρ, E) is $p_s = \mathrm{Tr}[\mathsf{E}\,\mathcal{E}_s(\rho)]$.

Choi-PSD constraints and smoothness.

Let $J(\Phi)$ be the Choi matrix of a map Φ. We fit $\{\mathcal{G}_i^{\mathrm{phys}}\}$, ρ, E, and (optionally) $(A, \{\mathcal{K}_r\})$ by minimizing a convex composite loss

$$\min \sum_{s\in\mathcal{S}} \Big\|\hat{p}_s - \mathrm{Tr}[\mathsf{E}\,\mathcal{E}_s(\rho)]\Big\|^2_{\Sigma_s^{-1}} + \lambda_{\mathrm{sm}} \sum_i \Big\|\mathcal{L}_{i+1}^{(0)} - \mathcal{L}_i^{(0)}\Big\|_F^2 \quad \text{s.t.} \quad \begin{cases} J(\mathcal{G}_i^{\mathrm{phys}}) \succeq 0, \quad \mathrm{Tr}_2\, J(\mathcal{G}_i^{\mathrm{phys}}) = \mathbb{I}, \\ J(\exp(\Delta t\, \mathcal{L}_i^{(0)})) \succeq 0 \quad \forall\, \Delta t \in \text{grid}, \end{cases} \tag{13.6}$$

plus linear constraints encoding gauge-fixing and instrument invariances. The penalty with weight λ_{sm} regularizes slow colored drift.

Goodness-of-fit and gates.

Report Pearson (or deviance) χ^2 per dof, and *CP margins*:

$$m_{\mathrm{CP}} = \min_i \lambda_{\min}\Big(J(\mathcal{G}_i^{\mathrm{phys}})\Big), \quad m_{\mathrm{TP}} = \min_i \Big\| \mathrm{Tr}_2\, J(\mathcal{G}_i^{\mathrm{phys}}) - \mathbb{I}\Big\|_\infty,$$

and (if embedding fitted) the *gap* $1-\rho(A)$. Acceptance on the DoV: $m_{\mathrm{CP}} \geq 0$, $m_{\mathrm{TP}} \leq \epsilon_{\mathrm{TP}}$, $\rho(A) < 1 - \epsilon_A$.

RB with colored noise: filter-function view.

For a random length-m Clifford sequence with interleaved idle or target gate, the survival probability obeys, in the colored-noise cumulant expansion,

$$\bar{F}(m) \;\approx\; A + B\,\exp\Big(-\chi(m)\Big), \qquad \chi(m) \;=\; \frac{1}{2\pi}\int_0^\infty \mathrm{d}\omega\; \mathsf{S}(\omega)\,\mathcal{F}_m(\omega), \tag{13.7}$$

where $\mathcal{F}_m(\omega)$ is the RB *sequence* filter derived from the toggle-frame two-point correlation of the randomized control sequence. *Markovian* noise gives $\chi(m) = m/m_\star$ (pure exponential); *quasi-static* (low-frequency dominated) noise yields $\chi(m) \approx \alpha\, m^2$ (Gaussian-like decay). For $1/f^\alpha$ spectra, $\chi(m) \propto m^{1+\alpha}$ up to a cutoff set by the gate rate.

Interleaved RB (gate-resolved).

Interleaving a gate G changes the filter to $\mathcal{F}_m^{(G)}$, defining an *effective error rate* $r_G \approx 1 - \exp(-\Delta\chi_G)$, $\Delta\chi_G = \chi^{(G)}(m) - \chi^{\mathrm{ref}}(m)$. Comparing $\Delta\chi_G$ to the prediction from the designed $Y(\omega)$ (via §12.1) cross-validates the compilation.

Leakage and SPAM handling.

Extend the RB model with a leakage term $C(-1)^m$ or a second exponential; fit SPAM parameters (A, B, C) jointly with $\chi(m)$. In GST, SPAM enters through ρ, E; we propagate their priors from the ledger (§12.3.2).

Cross-checks (GST ↔ RB ↔ spectroscopy).

From GST we extract a (time-averaged) GKSL generator $\overline{\mathcal{L}}$, predict an RB decay constant by linearizing $\chi(m)$ at the RB cycle rate, and compare with the RB fit; independently, the measured $S(\omega)$ from spectroscopy (Chap. 12) predicts $\chi(m)$ via (13.7). Agreement within envelopes is an acceptance criterion.

13.2.2 Error budgets

Additive variance decomposition.

Let θ be any scalar performance or physics parameter (e.g. a GKSL rate, a notch depth, or a coupling). For small perturbations δx_j with sensitivities $g_j = \partial\theta/\partial x_j$,

$$\mathrm{Var}(\widehat{\theta}) \simeq \underbrace{\mathrm{Var}_{\mathrm{stat}}}_{\text{shot/fit}} + \sum_j g_j^2\,\mathrm{Var}(\delta x_j) + \sum_{j\neq k} g_j g_k\,\mathrm{Cov}(\delta x_j, \delta x_k) + \underbrace{\mathrm{Var}_{\mathrm{model}}}_{\text{embedding/PSD misspec.}} + \underbrace{\mathrm{Var}_{\mathrm{DI}}}_{\text{colored/cone slacks}}.$$

We report a stacked budget with entries below.

Core budget entries (symbols, propagation).

- **SPAM** ($\epsilon_{\mathrm{SP}}, \epsilon_{\mathrm{M}}$): enter GST likelihood and RB offsets (A, B); propagate via the Fisher block $\Pi = \Sigma_\nu^{-1}$ in (12.20).
- **Timing jitter** (σ_t): contributes a derivative-kernel term to Q_{rob} (§12.1.2); in RB/GST adds an ω^2-weighted leakage in $\chi(m)$.
- **Pulse area** (δ_π): modifies the lab window H_τ and enters linearly in $\partial Y/\partial\tau_\pi$.
- **Detuning drift** (D_Δ): OU-process prior; induces correlated errors across sequence lengths; included as off-diagonals in Σ_d.
- **PSD modeling** (δS): difference between inverted $\widehat{S}(\omega)$ and parametric $S(\omega; u)$; propagated to $\chi(m)$ via the integral in (13.7).
- **Embedding bias** (q pseudomodes): model-order error for $(A, \mathcal{K}_r)$; evaluated by information criteria and residual whiteness tests (should be flat on DoV).
- **DI/cone slacks** ($\zeta_{\mathrm{col}}, \epsilon_{\mathrm{cone}}$): additive uncertainty on rates/amplitudes as specified in §7.3 and §11.1.2.

Budget table (reporting template).

For each θ we publish central value, statistical error, and systematic components with signs and correlations. A compact `tabularx` format:

Source	$\boldsymbol{\delta\theta}$	$\boldsymbol{\sigma}$	**Type**	**Notes**
Statistics (fit)	0	σ_{stat}	stat	From (12.19)
SPAM	$\partial_\theta \epsilon_{\text{SP}}\, \mu_{\text{SP}} + \partial_\theta \epsilon_{\text{M}}\, \mu_{\text{M}}$	$\sqrt{g^\top \Sigma_{\text{SPAM}} g}$	syst	Priors from calibration
Timing jitter σ_t	0	$\propto \lVert Q^{(\partial)} \rVert^{1/2}$	syst	See (12.11)
Pulse area δ_π	0	via $\partial Y / \partial \tau_\pi$	syst	Window H_τ
Detuning drift D_Δ	0	from OU fit	syst	Correlated in Σ_d
PSD model δS	bias	envelope integral	model	$\int \delta S\, \mathcal{F}$
Embedding order	bias	IC/whiteness	model	Pseudomode count q
DI/cone slacks	0	$\zeta_{\text{col}}, \epsilon_{\text{cone}}$	gate	Declared DoVslack
Total	$\delta\theta_{\text{tot}}$	σ_{tot}		Quadrature + correlations

Acceptance and stopping hooks.

Budgets must satisfy: (i) total systematic $\leq$ target fraction of statistical ($\sigma_{\text{syst}} \leq \kappa\, \sigma_{\text{stat}}$); (ii) DI/cone slacks within declared envelopes; (iii) residuals whitened on DoV. Violations trigger the stopping/revision rules in §12.3.2.

Artifact 13.2: AR: Verification packs

Deliver: (1) GST fit with Choi spectra, CP/TP margins and (if used) embedding $(A, \mathcal{K}_r)$; (2) RB decay data and filter-based $\chi(m)$ overlays; (3) inverted $\widehat{S}(\omega)$ with parametric overlays; (4) the full error budget tables and covariance objects; (5) gate compliance logs for $\mathsf{S}, \mathsf{TQ}, \mathsf{G}$.

13.3 Bridging to MSP Predictions: Maps to (S, T, U), EDM Floors, NSI Projections; Release Artifacts

Prelude. This section closes the loop between tabletop control data (filters, spectra, GKSL fits) and *phenomenological* MSP predictions used in §10–§11. We build explicit, unit-consistent maps from spectral/geometry parameters to: (i) oblique electroweak parameters (S, T, U); (ii) irreducible *EDM floors*; and (iii) neutrino NSI projections relevant to long-baseline appearance channels. We also specify reproducible *release artifacts* that encode these maps, their uncertainties, and gate certificates on the declared DoV.

13.3.1 Bridge maps

Chain-of-maps (overview).

$$\boxed{\text{Sequences} \Rightarrow Y(\omega) \Rightarrow \widehat{\mathsf{S}}(\omega),\ \widehat{\mathcal{L}} \Rightarrow \Theta_{\text{spec}} \Rightarrow \mathsf{G},\ \text{overlaps},\ C^{(d)} \Rightarrow (S,T,U),\ d_f,\ \varepsilon_{\alpha\beta}}$$

Here $\widehat{\mathsf{S}}(\omega)$ and $\widehat{\mathcal{L}}$ are, respectively, the inverted noise PSD and time-local GKSL generator (Chap. 12, 13.2); Θ_{spec} are pseudomode parameters $\{(\Omega_\ell, \Gamma_\ell, \lambda_\ell)\}$ and colored embedding coefficients; G is the fiber metric/mixing kernel; $C^{(d)}$ are EFT Wilson coefficients at μ_{ref} (from dispersion/threshold matching).

Källén–Lehmann bridge for gauge two-point functions.

Let $\Pi_{XY}(q^2)$ denote vacuum polarizations of electroweak currents ($X, Y \in \{3, B, 1\}$ in the weak basis). From the MSP spectral medium,

$$\Pi_{XY}(q^2) \;=\; \frac{1}{\pi}\int_0^\infty \frac{\rho_{XY}(s)}{s - q^2 - i0}\,\mathrm{d}s, \qquad \rho_{XY}(s) \;\geq 0 \;\; (\text{Gate } \mathsf{S}), \tag{13.8}$$

with ρ_{XY} computed from the colored embedding (pseudomodes) and fiber currents (Noether charges on the bundle). Low-energy EFT coefficients $C^{(6,8)}$ fix the subtractions. Derivatives at $q^2 = 0$ are

$$\Pi'_{XY}(0) \;=\; \frac{1}{\pi}\int_0^\infty \frac{\rho_{XY}(s)}{s^2}\,\mathrm{d}s, \qquad \Pi_{XY}(0) \;=\; \frac{1}{\pi}\int_0^\infty \frac{\rho_{XY}(s)}{s}\,\mathrm{d}s \;+\; \text{subtractions}. \tag{13.9}$$

Oblique parameters (S, T, U).

Using the standard definitions on the DoV,

$$\alpha S \;=\; 4 s_W c_W\, \Pi'_{3B}(0), \qquad \alpha T \;=\; \frac{\Pi_{11}(0) - \Pi_{33}(0)}{M_W^2}, \qquad \alpha U \;=\; 4 s_W^2 \Big[\Pi'_{11}(0) - \Pi'_{33}(0)\Big], \tag{13.10}$$

with $s_W = \sin\theta_W$, $c_W = \cos\theta_W$, and α the EM fine-structure constant. Combining (13.9) and threshold-matched $C^{(6)}$ yields

$$S \;=\; \frac{16\pi s_W c_W}{\pi\,\alpha}\int_0^\infty \frac{\rho_{3B}(s)}{s^2}\,\mathrm{d}s + S_{\text{sub}}[C^{(6)}], \;\; T \;=\; \frac{1}{\alpha M_W^2}\frac{1}{\pi}\int_0^\infty \frac{\rho_{11}(s) - \rho_{33}(s)}{s}\,\mathrm{d}s + T_{\text{sub}}[C^{(6)}], \tag{13.11}$$

(and similarly for U). The subtraction pieces $S_{\text{sub}}, T_{\text{sub}}$ are fixed by the Wilson set at μ_{ref} from Chap. 9.

CP-odd bridge and EDM floors.

Let $\mathcal{I}_{\text{CP}}$ be the leading CP-odd invariant generated by the MSP overlap geometry (e.g. a rephasing-invariant product of Yukawa-like overlaps with a geometric phase). The

dominant two-loop Barr–Zee–type floor for a fermion f (charge Q_f, mass m_f) is

$$d_f^{\text{floor}} \simeq \frac{e\, Q_f\, \alpha}{8\pi^3}\, m_f \sum_i \frac{\kappa_i \sin\phi_i}{M_i^2}\, J(\tau_i), \qquad \tau_i = \frac{m_i^2}{M_i^2}, \tag{13.12}$$

where M_i are mediator masses (gauge/pseudomode thresholds consistent with the DoV), κ_i are geometric couplings fixed by G and overlaps, and $J(\tau)$ a known loop kernel ($J > 0$ and $\mathcal{O}(1)$ on the DoV). The *floor* uses the *minimal* set $\{\kappa_i, \phi_i\}$ implied by MSP symmetries, hence cannot be tuned away without leaving the model class (§10.2). Uncertainties from hadronic/nuclear matrix elements enter as multiplicative envelopes in the kill-window test (11.25).

NSI projections.

Tabletop constraints map to effective four-fermion operators $\mathcal{L}_{\text{NSI}} = -2\sqrt{2}G_F\, \varepsilon_{\alpha\beta}^{fP}\, (\bar{\nu}_\alpha \gamma_\mu P_L \nu_\beta)(\bar{f}\gamma^\mu P f)$, with

$$\varepsilon_{\alpha\beta}^{fP} = \frac{1}{2\sqrt{2}G_F} \frac{g_{\alpha\beta}^{(fP)}}{\Lambda_{\text{eff}}^2}, \qquad g_{\alpha\beta}^{(fP)} = \sum_r \left(\mathcal{C}_r^{(6)}\right)_{\alpha\beta}^{fP}, \tag{13.13}$$

where $\mathcal{C}_r^{(6)}$ are the matched Wilsons from the MSP geometry (Lie–metric, abelian mixing, overlap textures; Chap. 9, 9.3). Device-independent bounds from §7.3 and tabletop filter certificates constrain combinations of $\varepsilon_{\alpha\beta}$; the bridge provides their propagation to appearance probabilities via the standard matter Hamiltonian $H_{\text{mat}} = \sqrt{2}G_F N_e(1+\varepsilon)$.

Uncertainty propagation and gates.

Let $\theta = (\Theta_{\text{spec}}, \mathsf{G}, \text{overlaps}, C^{(d)})$ be the bridge inputs with covariance Σ_θ (from verification packs). Then for any observable vector $O = (S, T, U, d_f, \varepsilon)$ with Jacobian $J_O = \partial O/\partial\theta$,

$$\Sigma_O \simeq J_O\, \Sigma_\theta\, J_O^\top, \qquad \text{acceptance: } O \in \mathcal{W}_p \text{ (§11.3.1) and gates } \mathsf{S}/\mathsf{G}/\mathsf{TQ}. \tag{13.14}$$

Dispersion prerequisites (analyticity, thresholds) are checked as in §11.1.1; cone LMIs for waveguide propagation as in §11.1.2.

13.3.2 Release artifacts

What we publish (per release tag).

- **Bridge cards (TeX/HTML).** Compact summaries of inputs θ, maps (13.10)–(13.13), central values and covariances (13.14), and pass/fail against decision windows.
- **HDF5 seeds.** Datasets for $\{\widehat{\mathsf{S}}(\omega),\ \widehat{\mathcal{L}},\ \Theta_{\text{spec}},\ \mathsf{G},\ C^{(d)},\ \text{overlaps}\}$, plus kernels used in integrals (spectral densities ρ_{XY}, subtraction data).
- **RG notebooks.** Scripts that transport couplings from μ_{ref} to M_Z (for $\sin^2\theta_W$ & α_3) with threshold sets and envelopes.

- **Cone/positivity certificates.** SDP logs and minimal eigenvalue margins for block LMIs and forward-limit LMIs (CSV + human-readable PDFs).
- **NSI/appearance pack.** JSON matrices of $\varepsilon_{\alpha\beta}^{fP}$ with priors/covariances; functions to evaluate modified $P_{\alpha\beta}(L, E)$ and orthogonal template overlays (§11.2.2).
- **EDM floor worksheet.** The CP-invariant $\mathcal{I}_{\mathrm{CP}}$, mediator set $\{M_i, \kappa_i, \phi_i\}$, loop kernel $J(\tau)$ and the propagated floor with hadronic envelope.
- **Provenance & CI.** Versioned hashes of all inputs/outputs; continuous integration tests that rebuild all bridge numbers and assert gate compliance.

Acceptance footer.

A release is *green* on the DoViff: (i) all bridge maps compile with nonnegative margins for Gate S/G/TQ; (ii) (S, T, U), $\sin^2\theta_W$, α_3 lie within their kill windows; (iii) EDM floors are above current experimental upper limits *only* within declared hadronic envelopes; (iv) NSI projections and ND–FD maps remain jointly feasible (§11.2.1).

Artifact 13.3: AR: Bridge & release bundle

A single archive containing bridge cards, HDF5 seeds, RG notebooks, gate certificates, NSI/appearance pack, EDM worksheet, and CI scripts. The bundle reproduces all numbers in this section from raw sequence packs and verification artifacts, with a one-command build and checksums.

Epilogue: Unified Closure Statement

What Has Been Proved and What Can Kill It

Prelude. On the declared DoV, Vol. V compresses the axioms and constructions of Vols. I–IV into a *master spectral* framework (MSP) with three pillars: (i) well-posed kinematics (cone/positivity), (ii) operational quantum dynamics (GKSL embeddings with colored noise and locality), and (iii) gauge/geometry consequences (Ward/BRST and coupling relations). This epilogue records, in a compact and testable form, *what is proved* and *what observations would invalidate it.* Each statement below inherits its explicit hypotheses and gates from the cited chapters; acceptance is always restricted to the DoV.

Closure Set

Theorems (proved on the DoV).

1. **Born-limit (ergodic flooding).** Under (H1–H5), with a CP time-local (possibly colored) embedding and a Liouvillian mixing gap $\Delta_{\mathcal{L}} > 0$ on pointer sectors (§3), the channel admits a unique stationary state ρ_∞ and for any calibrated pointer POVM $\{\Pi_i\}$:

$$\lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \mathbf{1}\{X_t = i\} \;=\; p_i \;=\; \mathrm{Tr}(\Pi_i\, \rho_\infty) \qquad \text{a.s.} \tag{13.15}$$

   Finite-time deviations obey sub-Gaussian bounds controlled by the locking rate γ (§3.3).

2. **No-signalling (local GKSL).** For spacelike-separated instruments A and B with supports separated by the locality radius L_{loc} on (M, g_{eff}) (§4), the reduced statistics of A are independent of the choice of B:

$$\mathrm{Tr}_B\Big[\mathcal{E}_{AB}(\rho)\Big] \;=\; \mathcal{E}_A\Big(\mathrm{Tr}_B\, \rho\Big), \qquad \forall\, \rho, \;\; \forall\, \text{CP instruments } \mathsf{B} \text{ on } B. \tag{13.16}$$

   The result extends to colored noise via CP time-local embeddings (§4.2.2).

3. **Ward/BRST from MSP.** From the MSP minimal axioms S0–S4 (§8) and the

emergent Ehresmann→YM connection (§8.2), the Slavnov–Taylor (Ward) identities hold on the physical quotient:

$$\mathcal{S}(\Gamma) = 0, \qquad H^0_{\mathrm{BRST}}(\text{fields}) \simeq \text{observables}, \quad \text{anomalies controlled by inflow terms}, \tag{13.17}$$

with cohomological equivalence between BRST-closed modulo exact functionals and gauge-invariant observables (§8.3).

Predictions (numerical/structural).

1. **Gauge-coupling relations.** The Lie–metric on $\mathrm{Iso}(\mathcal{F})$ fixes algebraic relations among (g_1, g_2, g_3) at a reference scale μ_{ref}, yielding an allowed ellipsoid in $(\alpha_1, \alpha_2, \alpha_3)$ (§9.1).
2. **Weinberg angle.** A *reference-scale* prediction for $\sin^2\theta_W(\mu_{\mathrm{ref}})$ with uncertainties propagated to M_Z via RG/threshold maps (§9.2.1, §9.2.2).
3. **Textures and sum rules.** Overlap-driven Yukawa textures imply CKM/PMNS sum rules and phase relations, testable against global fits on the DoV(§9.3).

Acceptance note.

Each theorem/prediction is paired with a *kill window* and gate checks in Chap. 11: positivity/hyperbolicity (cone LMIs), dispersion forward limits, device-independent feasibility, and ND–FD consistency. Failure of any earlier kinematic gate (S/G) dominates later, model-dependent windows, and triggers the revision playbook (§11.3.2).

Kill Signatures

Principle. A *kill signature* is recorded when duly calibrated observations, restricted to the declared DoVand passing instrument checks (H5), fall *outside* the acceptance window of a theorem or prediction while all earlier kinematic gates (S/G) are enforced. Each window below is defined by an explicit statistic/inequality; failure triggers the revision playbook (§11.3.2) with the *dominance order*

$$\text{Cone/positivity} \Rightarrow \text{Forward-limit} \Rightarrow \text{ND–FD} \Rightarrow \text{Precision/EDM} \Rightarrow \text{Gauge/NSI}.$$

1. **Positivity & hyperbolicity (cone/forward).**
   - *Cone LMIs (gravity/causality).* Let $\mu_{\min}$ be the minimal eigenvalue over sampled directions of the block matrix in §11.1.2. **Kill if** $\mu_{\min} < 0$ (beyond numerical tolerance).
   - *Forward-limit positivity (dispersion).* For each tested channel, let $\lambda^{(2k)}_{\min}$ be the smallest eigenvalue of the Gram–projected Wilson blocks $\mathcal{C}^{(2k)}$ (§11.1.1). **Kill**

if $\lambda_{\min}^{(2k)} < 0$ for any k on the DoV.

2. **Oblique electroweak parameters** (S, T, U)**.** Let $O = (S, T, U)$ with MSP prediction $O_\star$ and total covariance Σ_O (including thresholds/RG and declared slacks; §10.2). Define $\chi_O^2 = (\widehat{O} - O_\star)^\top \Sigma_O^{-1} (\widehat{O} - O_\star)$. **Kill if** $\chi_O^2 > \chi_{p,3}^2$ (outside the p-level ellipsoid) on the DoV.

3. **EDM floors (irreducible CP).** For fermion f, the MSP floor d_f^{floor} from (13.12) with hadronic/nuclear envelope δ_{had} implies the one-sided test

$$|d_f|_{\text{exp}}^{95\%} \;<\; d_f^{\text{floor}}\,(1 - \delta_{\text{had}}) \quad \Longrightarrow \quad \textbf{Kill (CP floor violated)}.$$

   Applies to $d_e, d_n, \ldots$ with platform-appropriate matrix-element priors.

4. **Device-independent (DI) bounds under colored noise.** Using the DI feasibility regions constructed with colored slacks (§7.3), let $\widehat{\mathbf{c}}$ be the vector of observed correlators (e.g. CHSH and ancillary moments). **Kill if** $\widehat{\mathbf{c}}$ lies outside the certified DI region compatible with the CP time-local embedding (even after allowed slacks) or if a no-signalling linear functional exceeds its bound (violating §4.2).

5. **Near–far (ND–FD) consistency.** Construct the ND–FD transfer $\widehat{n}^F$ from ND data as in (11.15), residual $\mathbf{r} = n^F - \widehat{n}^F$, and covariance Σ_{NF} (§11.2.1). **Kill if** $p_{NF} = \Pr(\chi^2 \geq \mathbf{r}^\top \Sigma_{NF}^{-1} \mathbf{r})$ falls below the gate threshold p_{gate} on the DoV.

Audit on kill.

Every kill event must log: (i) the window statistic value and threshold; (ii) the active DoVand instrument calibration hash; (iii) gate certificates (S/G/TQ) at the time of test; (iv) the automatic branch taken in the revision playbook (§11.3.2) and the smallest change needed to recover feasibility (if any). All logs are shipped as part of the public release artifacts (§13.3.2).

Artifacts & Reproducibility

Policy. Every numerical statement and gate decision in Vol. V must be reproducible from publicly released artifacts with a *one-command* rebuild on the declared DoV. We require deterministic seeds, pinned environments, provenance manifests, and continuous–integration (CI) tests that assert all acceptance gates (S/G/TQ) and decision windows pass within published margins.

Must-ship bundle (per release tag).

The public bundle contains the following top-level items; each has machine- and human-readable forms and is versioned by cryptographic hash.

1. **Cards / UFO (Unified Field Overview).** A compact PDF/HTML card-set summarizing: (i) theorems proved (Born-limit, no-signalling, Ward/BRST) and their hypotheses; (ii) master predictions ($g_{1,2,3}$ relations, $\sin^2\theta_W$, textures/sum rules); (iii) decision windows and pass/fail with margins on the DoV; (iv) links to certificates and raw data. Each card is auto-generated from the HDF5 seeds and CI logs.

2. **RG notebooks.** Executable notebooks (and a headless script) that:

   - load reference-scale couplings and threshold sets;
   - transport to M_Z with envelopes;
   - emit tables/plots for $\sin^2\theta_W(M_Z)$, $\alpha_3(M_Z)$ and (S, T, U) correlations;
   - write back results into the HDF5 bundle (with checksum).

3. **HDF5 seeds.** A single file `msp_release.h5` with the schema:

```
/provenance/manifest.yml        # commit hashes, env, seeds, DoV
/filters/Y_omega                # |Y(ω)|^2 arrays per sequence
/filters/y_t                    # time-domain controls per sequence
/kernels/Q, R, B_l, Q_partial   # design & robust kernels
/verification/GKSL/choi         # Choi spectra, CP/TP margins
/verification/RB                # m, F(m), fit params, χ(m)
/spectra/S_omega                # inverted PSD with covariances
/couplings/alphas_ref           # (α1, α2, α3) at μ_ref
/couplings/RG_to_MZ             # transported couplings + cov
/oblique/STU                    # (S,T,U) + covariances
/edm/floors                     # d_f^floor with envelopes
/nsi/epsilon                    # ϵ_{αβ}^{fP} matrices + priors
/windows/decision               # statistics, thresholds, pass/fail
/gates/cone, forward            # LMI certificates, min eigenvalues
/ci/logs                        # test outputs, timings, resources
```

4. **Filter library.** A catalog of reference sequences (timings or tone sets) with certified notch metrics and robustness margins:

   - `sequences/`*name*`.json` (hardware-neutral);
   - export adapters for cold atoms (AOM/EOM) and waveguides (IQ).

5. **CI regression tests.** A scripted pipeline that rebuilds the full stack from `msp_release.h5` and asserts:

   (a) *Cone/positivity:* minimal LMI eigenvalues ≥ 0 (with tolerances);

   (b) *Forward-limit:* Gram–projected Wilson blocks $\succeq 0$;

(c) *Verification:* GST CP/TP margins nonnegative; RB decay within filter-based envelopes;

(d) *Predictions:* RG numbers match stored values within stated numerical tolerances;

(e) *Windows:* all decision statistics inside acceptance unless explicitly flagged as `kill`;

(f) *Determinism:* re-running with the same seeds yields bitwise-identical outputs (where applicable).

Provenance & environment pinning.

Each bundle includes:

- a `manifest.yml` with code hashes, dataset checksums, random seeds, platform info, and the DoVdeclaration;
- an environment lockfile (e.g. `conda-lock.yml` or `requirements.txt` plus compiler/runtime versions);
- a *container recipe* (Docker/Singularity) to guarantee portability.

Reproduction quickstart (one-command build).

1. Create the environment from the lockfile or container.
2. Run `make release_replay` (or the provided `build.py`): this regenerates RG numbers, verifies LMI certificates, replays GST/RB fits, and rebuilds the Cards/UFO from `msp_release.h5`.
3. Inspect CI report; all gates must pass on the DoVwith published margins.

Versioning, DOIs, and licensing.

Releases are semantic-versioned; HDF5 seeds and Cards/UFO receive persistent identifiers (DOIs). Code is licensed for reproducible research; data artifacts (Cards/UFO, seeds) carry a permissive data license with attribution.

Artifact 13.4: Public bundle and CI badges

We publish (i) the `msp_release.h5` seeds, (ii) Cards/UFO, (iii) RG notebooks and headless scripts, (iv) filter library with adapters, (v) cone/forward certificates, and (vi) CI logs and badges indicating gate status on the declared DoV.

Conceptual Afterword: Event-Waves, Spectral Flooding, and Emergence

Why a conceptual afterword? Vol. V established an operational closure: concrete gates, hypotheses, and falsifiable windows on a declared DoV. This afterword distills the conceptual stance behind those constructions—clarifying what "exists" in the theory, what is merely instrument-relative, and how the Master Spectral Principle (MSP) organizes emergence and prediction—so that Vol. VI can focus on audit and procedural details without re-motivating the philosophy.

From Ontology to Operation

Spectral medium, not a mechanical ether.

The unit-bearing descriptors $(\rho, \eta, \phi, T, \chi^\mu)$ (H0) define a *spectral* medium: a bookkeeping of intensities, couplings, and transport cones that are *operationally* accessible on the DoV. It is neither a luminiferous ether nor a hidden-variable substrate; rather, it is a structured set of constraints that (i) fix causal cones (g_{eff}), (ii) bound admissible instruments (H5), and (iii) seed gauge geometry (MSP).

Event-waves vs. state-waves.

An *event-wave* is a propagating constraint on *instrumental events* (counts, clicks, pointer deflections), not a literal substance. Mathematically it is the transport of sensitivity functionals through the cone determined by g_{eff}; physically it tells us which joint statistics are even *possible* given locality (H1) and CP embeddings. "State-waves" (wavefunctions) are computational devices in Hilbert space; event-waves are feasibility flows in the space of observations.

Emergence is a contract, not a miracle.

g_{eff} and Ehresmann→YM connections emerge as *coarse-grained contracts*: once pointer sectors stabilize (gaps, §2.2) and Noether currents are fibered (§2.3), the geometry required to consistently quotient redundancies is *forced*. MSP states this as a minimal axiom set with Ward/BRST as consequences (§8, §8.3).

Spectral Flooding and the Born-limit

Flooding as ergodic accounting.

"Spectral flooding" names the relaxation of energy/share ledgers under local, CP-divisible dynamics into pointer-stationary weights. The Born-limit theorem in §3 identifies the asymptotic frequencies with $p_i = \mathrm{Tr}(\Pi_i \rho_\infty)$ once (i) mixing gaps persist, (ii) channels are calibrated (H5), and (iii) locality forbids superluminal coordination (§4). Thus the probabilistic postulate is traded for a dynamical limit with deviation bounds (§3.3).

Zeno engineering, not fine-tuning.

Colored-Zeno control (§5.2) *shapes* which bands the medium can flood through, but cannot create signalling or violate cones; convex LMIs encapsulate those non-negotiable constraints (§5.2.2, §11.1.2). Control is epistemic leverage inside the gates.

Threads, Nonseparability, and No-signalling

Threads as structural resources.

χ^μ-threads quantify nonseparability constrained by the cone: they are not wires carrying influences faster than g_{eff} permits. Thread-adjusted measures and DI regions translate that into certifiable bounds (§6, §7.3).

Tsirelson without superluminal shortcuts.

Local GKSL $\Rightarrow$ no-signalling (§4.2) survives colored embeddings (§4.2.2). Maximal CHSH is attainable only within those embeddings; DI slack accounts for colored statistics (§7).

MSP, Minimal Inputs, and Prediction

Minimal sovereign inputs.

MSP compresses Vols. I–IV into a handful of inputs (e.g. $\kappa_G, \mathcal{N}_u, \mathcal{N}_d, \mathcal{N}_\ell$ and topological indices), from which geometry fixes coupling relations and textures (§9, §9.3). This is *structural realism*: what survives model changes are relations (cones, symmetries, overlaps), not microscopic stories.

Falsifiability by windows, not anecdotes.

Decision tables (§11.3) tie each claim to a kill window: cone/forward positivity, (S, T, U), EDM floors, DI bounds, and ND–FD maps. This enforces Popperian discipline while acknowledging DoV-limited modeling.

Epistemic Discipline: The DoVContract

Why DoVfirst.

Every statement in Vol. V is *conditional* on a declared DoV. The boundary triggers (gap closure, loss of CP-divisibility, locality violation, negative LMI eigenvalues) are not merely technicalities—they are the *rules of the game* that prevent category errors (e.g. extrapolating beyond validated scales). Vol. VI turns this into an audit protocol (instrument checks, CI, seeds).

Glossary of Core Notions (one-liners)

- **Event-wave:** cone-respecting propagation of feasibility constraints on instrument events.
- **Spectral flooding:** ergodic equilibration of pointer energy/share ledgers $\Rightarrow$ Born-limit.
- **Emergent g_{eff}:** operational cone/metric needed to make quotients and local GKSL well-posed.
- **Thread:** cone-constrained nonseparability resource; quantified by thread-adjusted measures.
- **MSP:** minimal axioms enforcing spectral/geometry relations; yields Ward/BRST and coupling constraints.
- **DoV:** maximal region where gates hold and embeddings/mixing gaps are certified.

Looking Ahead to Vol. VI

From concept to audit.

Vol. VI operationalizes this stance: (i) formal DoVaudits and boundary scans; (ii) instrument certification and SPAM ledgers; (iii) CI and artifact policies; (iv) reproducible inversion of $\widehat{\mathsf{S}}(\omega)$ and GKSL fits; and (v) governance of gate failures (revision playbook).

Promise.

If Vol. V says *what counts as proof and what can kill it*, Vol. VI says *how you prove it every time, from raw counts to CI badges.*

Appendix: Concept–Mathematics Bridge

Purpose. This appendix gives a compact mapping from the conceptual vocabulary (event-waves, spectral flooding, emergence, threads, DoV, ...) to the precise mathematical structures and operational tests used throughout Vol. V. It is a reader's "legend" for cross-linking philosophy to proofs and kill windows.

Concept	Mathematical avatar	Operational gate/test	Primary refs
Event-wave	Transport of instrument–feasibility functionals along the emergent cone g_{eff}; hyperbolic PDE well-posedness; locality radius L_{loc}	Cone LMIs (gravity/causality), ND–FD map consistency, no-signalling checks	Sections 4.2, 11.1.2 and 11.2
Spectral flooding	Energy/share balance kernels; ergodic/martingale convergence to stationary pointer weights $p_i = \text{Tr}(\Pi_i \rho_\infty)$; locking rate γ	Born-limit theorem; finite-time sub-Gaussian deviation bounds vs. γ	Chapter 3 and Section 3.3
Emergence (g_{eff}, connection)	Pre-symplectic reduction → CCR/CAR; Ehresmann → Yang–Mills connection on fibers; physical-quotient cohomology	Forward-limit positivity; Ward/ST identities on physical quotient	**??** and Sections 8.2 and 8.3
Threads (χ^μ)	Cone-constrained nonseparability; thread-adjusted negativity/rel.-entropy; SDP witnesses	DI feasibility under colored noise; monogamy/cut inequalities	Chapter 6 and Sections 6.2 and 7.3
No-signalling	Local GKSL CP semigroups; CP time-local embeddings for colored noise	No-signalling theorem; DI linear-functionals bounded; channel tomography CP margins	Chapter 4 and Sections 4.2 and 13.2
Colored-Zeno control	Filter functions $\lvert Y(\omega)\rvert^2$; Toeplitz/SOS LMIs for trigonometric nonnegativity; convex SDP (Gram lifting)	Notch/band constraints satisfied; robustness margins $\{m_{\text{pow}}, m_{\text{band}}, m_\partial, m_J\}$	Sections 5.2 and 12.1 and **??**
DoV(Domain of Validity)	Intersection where: cone/forward LMIs hold; CP time-local embedding exists; mixing gap $\Delta_{\mathcal{L}} > 0$; calibrations within budget	Boundary triggers (gap closure, CP-loss, locality violation, negative eigenvalues)	Theorem 1.1 and **??**
Gauge-coupling geometry	Lie-metric on Iso($\mathcal{F}$); relations among (g_1, g_2, g_3); reference-scale ellipsoid	Ellipsoidal acceptance for $(\alpha_1, \alpha_2, \alpha_3)$; RG transport to M_Z	Sections 9.1 and 10.1 and **??**
Textures and sum rules	Overlap integrals on fibers ⇒ Yukawa textures; CKM/PMNS correlations	Global-fit consistency; sum-rule residuals within windows	Section 9.3
EDM floor (irreducible CP)	CP-odd invariant $\mathcal{I}_{\text{CP}}$; Barr–Zee–type two-loop floors with $\{\kappa_i, \phi_i\}$	One-sided kill: $\lvert d\rvert^{\text{UL}} < d^{\text{floor}}(1 - \delta_{\text{had}})$	Sections 10.2 and 13.3
NSI projections	Wilson $C^{(6)} \to \varepsilon^{fP}_{\alpha\beta}$; matter Hamiltonian with ε	Appearance-channel likelihoods; ND–FD orthogonal templates	Sections 10.3, 11.2 and 13.3

Gate legend (shorthand).

- **S** — Stability/positivity: Hadamard, CP, PSD/Toeplitz, forward-limit positivity.
- **G** — Geometry/cone: emergent g_{eff}, locality radius L_{loc}, cone LMIs.

- **TQ** — Tomographic quantum: CPTP channel consistency (GST/RB), calibrated instruments (H5).
- **D (Decision)** — Windows and kill tests: (S, T, U), EDM floors, DI regions, ND–FD maps.

Mini "dictionary".

- "*Wave of events*" $\leftrightarrow$ hyperbolic transport of feasible joint statistics constrained by g_{eff} (G), verified by ND–FD consistency.
- "*Flooding*" $\leftrightarrow$ martingale/ergodic convergence of pointer ledgers $\Rightarrow$ Born-limit with finite-time bounds (S).
- "*Emergence*" $\leftrightarrow$ necessity of Ehresmann$\to$YM structure to preserve quotient consistency and Ward/BRST.
- "*Threads*" $\leftrightarrow$ DI-certified, cone-limited nonseparability quantified by thread-adjusted measures.
- "DoV" $\leftrightarrow$ the audited region where all gates hold and boundary triggers are monitored.